

US010516890B2

(12) **United States Patent**
Wang et al.

(10) **Patent No.:** **US 10,516,890 B2**
(45) **Date of Patent:** **Dec. 24, 2019**

(54) **ACCELERATING MACHINE OPTIMISATION PROCESSES**

(71) Applicant: **Magic Pony Technology Limited**,
London (GB)

(72) Inventors: **Zehan Wang**, London (GB); **Robert David Bishop**, London (GB); **Wenzhe Shi**, London (GB); **Jose Caballero**, London (GB); **Andrew Peter Aitken**, London (GB); **Johannes Totz**, London (GB)

(73) Assignee: **Magic Pony Technology Limited**,
London (GB)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **15/680,708**

(22) Filed: **Aug. 18, 2017**

(65) **Prior Publication Data**
US 2018/0129918 A1 May 10, 2018

Related U.S. Application Data

(63) Continuation of application No. PCT/GB2016/050427, filed on Feb. 19, 2016.

(30) **Foreign Application Priority Data**

Feb. 19, 2015 (GB) 1502753.5
Feb. 27, 2015 (GB) 1503427.5
(Continued)

(51) **Int. Cl.**
G06T 3/00 (2006.01)
H04N 19/36 (2014.01)
(Continued)

(52) **U.S. Cl.**
CPC **H04N 19/36** (2014.11); **G06K 9/46** (2013.01); **G06K 9/6215** (2013.01); **G06K 9/66** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC .. H04N 19/117; H04N 19/142; H04N 19/154; H04N 19/172; H04N 19/177;
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,638,125 A 6/1997 Ahn, et al.
5,673,213 A 9/1997 Weigl
(Continued)

FOREIGN PATENT DOCUMENTS

CN 101621683 A 1/2010
CN 102148996 A 8/2011
(Continued)

OTHER PUBLICATIONS

U Dong Chao et al: "Learning a Deep Convolutional Network for Image Super-Resolution", Sep. 6, 2014 (Sep. 6, 2014), *Correct System Design*; Springer International Publishing, Cham pp. 184-199, XP047296566, ISSN:0302-9743, ISBN:978-3-642-36361-0 (Year: 2014).*

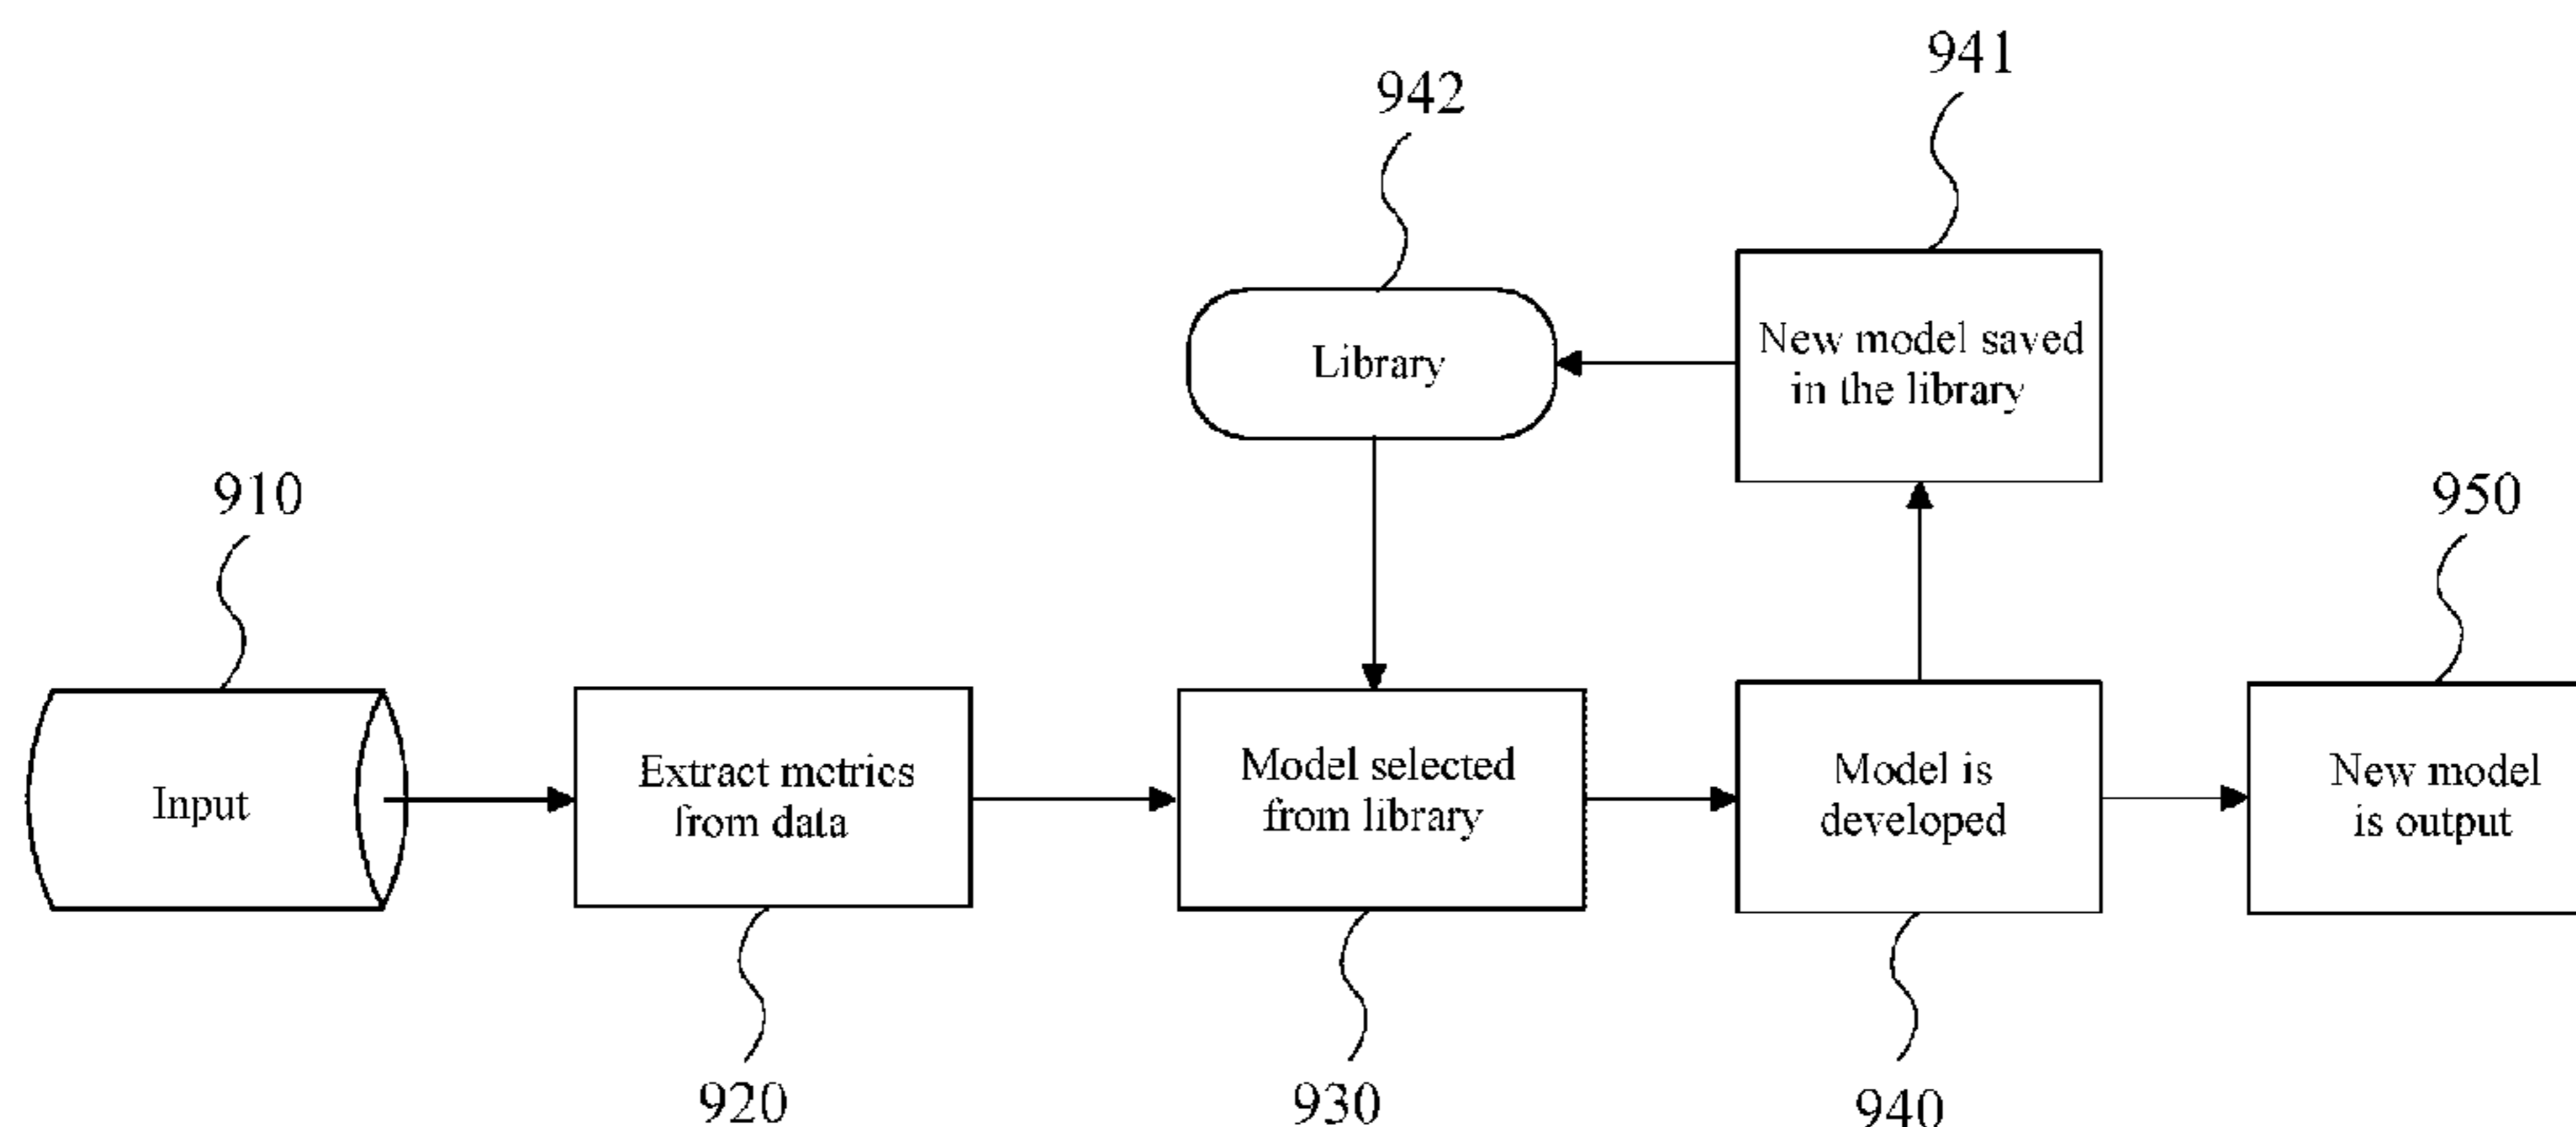
(Continued)

Primary Examiner — Ping Y Hsieh
(74) *Attorney, Agent, or Firm* — Brake Hughes Bellermann LLP

(57) **ABSTRACT**

A method for training learned hierarchical algorithms, the method comprising the steps of receiving input data and generating metrics from the input data. At least one hierarchical algorithm is then selected from a plurality of predetermined hierarchical algorithms based on comparing the generated metrics from the input data and like metrics for
(Continued)

900



each of the plurality of predetermined hierarchical algorithms. The selected hierarchical algorithm is developed based on the input data and the developed hierarchical algorithm is outputted.

20 Claims, 27 Drawing Sheets

G06T 5/001; G06T 5/002; G06T 2207/10016; G06T 2207/20081; G06T 2207/20084; G06K 9/46; G06K 9/6215; G06K 9/66; G06N 3/04; G06N 3/0445; G06N 3/049; G06N 3/08

See application file for complete search history.

(30) Foreign Application Priority Data

Mar. 21, 2015	(GB)	1505544.5
Apr. 27, 2015	(GB)	1507141.8
May 21, 2015	(GB)	1508742.2
Jun. 25, 2015	(GB)	1511231.1
Nov. 3, 2015	(GB)	1519425.1
Nov. 6, 2015	(GB)	1519687.6

(51) Int. Cl.

G06T 3/40	(2006.01)
H04N 19/46	(2014.01)
H04N 19/117	(2014.01)
H04N 19/142	(2014.01)
H04N 19/154	(2014.01)
H04N 19/177	(2014.01)
H04N 19/80	(2014.01)
H04N 19/86	(2014.01)
H04N 19/59	(2014.01)
G06T 5/00	(2006.01)
H04N 19/31	(2014.01)
H04N 19/33	(2014.01)
G06N 3/04	(2006.01)
G06N 3/08	(2006.01)
H04N 7/01	(2006.01)
G06T 7/11	(2017.01)
G06K 9/62	(2006.01)
G06K 9/66	(2006.01)
H04N 19/172	(2014.01)
G06K 9/46	(2006.01)
H04N 19/87	(2014.01)
H04N 19/176	(2014.01)

(52) U.S. Cl.

CPC **G06N 3/04** (2013.01); **G06N 3/049** (2013.01); **G06N 3/0445** (2013.01); **G06N 3/0454** (2013.01); **G06N 3/08** (2013.01); **G06T 3/40** (2013.01); **G06T 3/4007** (2013.01); **G06T 3/4046** (2013.01); **G06T 3/4053** (2013.01); **G06T 5/001** (2013.01); **G06T 5/002** (2013.01); **G06T 7/11** (2017.01); **H04N 7/0117** (2013.01); **H04N 19/117** (2014.11); **H04N 19/142** (2014.11); **H04N 19/154** (2014.11); **H04N 19/172** (2014.11); **H04N 19/177** (2014.11); **H04N 19/31** (2014.11); **H04N 19/33** (2014.11); **H04N 19/46** (2014.11); **H04N 19/59** (2014.11); **H04N 19/80** (2014.11); **H04N 19/86** (2014.11); **G06T 2207/10016** (2013.01); **G06T 2207/20081** (2013.01); **G06T 2207/20084** (2013.01); **H04N 19/176** (2014.11); **H04N 19/87** (2014.11)

(58) Field of Classification Search

CPC H04N 19/31; H04N 19/33; H04N 19/46; H04N 19/59; H04N 19/80; H04N 19/86; H04N 7/0117; H04N 19/176; H04N 19/87; G06T 7/11; G06T 3/40; G06T 3/4007; G06T 3/4046; G06T 3/4053;

(56)

References Cited

U.S. PATENT DOCUMENTS

5,831,678	A	11/1998	Proctor
6,075,884	A	6/2000	Lubin et al.
6,470,097	B1	10/2002	Lai et al.
6,735,336	B2	5/2004	Avni et al.
7,133,568	B2	11/2006	Nikitin et al.
7,477,800	B2	1/2009	Avidan et al.
7,623,731	B2	11/2009	Lim et al.
7,676,441	B2	3/2010	Matsugu et al.
8,037,010	B2	10/2011	Jaros et al.
8,311,967	B1	11/2012	Lin et al.
8,331,615	B2	12/2012	Furukawa et al.
8,467,599	B2	6/2013	El
8,731,255	B2	5/2014	El-Baz
9,208,539	B2	12/2015	Choudhury et al.
9,342,870	B2	5/2016	Lin et al.
9,824,421	B2	11/2017	Ragozin et al.
9,904,889	B2	2/2018	Eliasmith et al.
9,906,787	B2	2/2018	Amon et al.
10,019,642	B1	7/2018	Navarrete Michelini et al.
10,181,092	B2	1/2019	Shao et al.
2003/0071908	A1	4/2003	Sannoh et al.
2004/0086039	A1	5/2004	Reynolds et al.
2004/0136567	A1	7/2004	Billinghurst et al.
2004/0150538	A1	8/2004	Kim et al.
2004/0218834	A1	11/2004	Bishop et al.
2005/0063461	A1	3/2005	Lee et al.
2005/0200757	A1*	9/2005	Pica H04N 19/30 348/390.1
2007/0172452	A1	7/2007	Klucher et al.
2007/0223887	A1	9/2007	Kanamori et al.
2009/0034622	A1	2/2009	Huchet et al.
2009/0067491	A1	3/2009	Sun et al.
2009/0110285	A1	4/2009	Elad et al.
2009/0232213	A1	9/2009	Jia
2010/0150229	A1*	6/2010	Francois H04N 19/70 375/240.12
2010/0183217	A1	7/2010	Seung et al.
2011/0299770	A1	12/2011	Vaddadi et al.
2011/0304687	A1	12/2011	Joshi et al.
2012/0288015	A1	11/2012	Zhang et al.
2012/0294369	A1*	11/2012	Bhagavathy H04N 19/119 375/240.16
2013/0128111	A1	5/2013	Corral-Soto
2013/0223734	A1	8/2013	Tuzel et al.
2014/0177706	A1	6/2014	Fernandes et al.
2014/0204996	A1	7/2014	Spears et al.
2014/0301661	A1	10/2014	Voronov et al.
2014/0369401	A1	12/2014	Minoo et al.
2015/0015569	A1	1/2015	Jung et al.
2015/0086109	A1	3/2015	Paris et al.
2015/0181159	A1	6/2015	Matsuyama
2016/0078600	A1	3/2016	Perez Pellitero et al.
2016/0140408	A1	5/2016	Shen et al.
2017/0132760	A1	5/2017	Kothule et al.
2017/0193680	A1	7/2017	Zhang et al.
2017/0345130	A1	11/2017	Wang et al.
2017/0347060	A1	11/2017	Wang et al.
2017/0347061	A1	11/2017	Wang et al.
2017/0347110	A1	11/2017	Wang et al.
2017/0374374	A1	12/2017	Wang et al.
2018/0121769	A1	5/2018	Huszar et al.
2018/0122048	A1	5/2018	Wang et al.
2018/0130177	A1	5/2018	Wang et al.
2018/0130178	A1	5/2018	Wang et al.
2018/0130179	A1	5/2018	Wang et al.
2018/0130180	A1	5/2018	Wang et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

2018/0131953 A1 5/2018 Wang et al.
 2018/0197089 A1 7/2018 Krasser et al.
 2018/0293713 A1 10/2018 Vogels et al.

FOREIGN PATENT DOCUMENTS

EP 0618737 A2 10/1994
 EP 1720358 A2 11/2006
 EP 2557789 A2 2/2013
 EP 2806374 A1 11/2014
 JP 1021385 A 1/1998
 KR 20030009575 A 2/2003
 WO 9739417 A2 10/1997
 WO 0033207 A1 6/2000
 WO 02/005208 A3 1/2002
 WO 03094034 A1 11/2003
 WO 2008133951 A2 11/2008
 WO 2013143396 A1 10/2013
 WO 2014052740 A1 4/2014
 WO 2015193531 A1 12/2015
 WO 2016132145 A1 8/2016
 WO 2016132147 A1 8/2016

OTHER PUBLICATIONS

Cayton, "Algorithms for manifold learning", Jun. 15, 2005, 17 pages.

Chopra, et al., "Learning a similarity metric discriminatively, with application to face verification", In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2005, pp. 539-549.

Dong, et al., "Compression Artifacts Reduction by a Deep Convolutional Network", IEEE International Conference on Computer Vision, Apr. 27, 2015, pp. 576-584.

Dong, et al., "Learning a Deep Convolutional Network for Image Super-Resolution", Lecture Notes in Computer Science, Sep. 6, 2014, pp. 184-199.

Search Report for GB Application No. 1502753.5, dated Jan. 12, 2016, 4 pages.

Search Report for GB Application No. 1503427.5, dated Dec. 17, 2015, 4 pages.

Search Report for GB Application No. 1505544.5, dated Feb. 12, 2016, 9 pages.

Search Report for GB Application No. 1507141.8, dated Dec. 21, 2015, 4 pages.

Search Report for GB Application No. 1508742.2, dated Jan. 25, 2016, 5 pages.

Search Report for GB Application No. 1511231.1, dated Jan. 11, 2016, 5 pages.

Search Report for GB Application No. 1519687.6, dated Apr. 13, 2016, 5 pages.

Examination Report for GB Application No. 1617478.1, dated May 3, 2017, 4 pages.

Examination Report for GB Application No. 1617478.1, dated Nov. 18, 2016, 6 pages.

Kim, et al., "Example-based Learning for Single-Image Super-Resolution and JPEG Artifact Removal", Technical Report No. TR-173, retrieved on Jun. 1, 2016 from http://www.kyb.mpg.de/fileadmin/user_upload/files/publications/attachments/TechReport-173, Aug. 2008, 28 pages.

Liu, et al., "Nonparametric Scene Parsing: Label Transfer via Dense Scene Alignment", IEEE Conference on Computer Vision and Pattern Recognition, Jun. 2009, 8 pages.

Matikainen, et al., "Model Recommendation for Action Recognition", IEEE Conference on Computer Vision and Pattern Recognition, 2012, pp. 2256-2263.

Mittal, et al., "No-Reference Image Quality Assessment in the Spatial Domain", IEEE Transactions on Image Processing, vol. 21, Issue 12, Dec. 2012, pp. 4695-4708.

Pan, et al., "A Survey on Transfer Learning", IEEE Transactions on Knowledge and Data Engineering, vol. 22, No. 10, Oct. 2010, pp. 1345-1359.

International Search Report and Written Opinion for PCT Application No. PCT/GB2016/050423, dated May 11, 2016, 16 pages.

International Search Report and Written Opinion for PCT Application No. PCT/GB2016/050425, dated May 4, 2016, 18 pages.

Torralba, et al., "80 million tiny images: a large dataset for non-parametric object and scene recognition", IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 30, Issue 11, Nov. 2008, pp. 1958-1970.

Wagstaff, et al., "Constrained K-means clustering with background knowledge", Proceedings of the Eighteenth International Conference on Machine Learning, 2001, pp. 577-584.

Yang, et al., "Coupled Dictionary Training for Image Super-Resolution", IEEE Transactions on Image Processing, vol. 21, Issue 8, Aug. 2012, 27 pages.

Yang, "Image super-resolution via sparse representation", IEEE Transactions on Image Processing, vol. 19, Issue 11, Nov. 2010, 51 pages.

Zeyde, et al., "On Single Image Scale-Up Using Sparse-Representations", Curves and Surfaces, Springer, 2012, pp. 711-730.

CS231n, "Convolutional Neural Networks for Visual Recognition"; Feb. 16, 2015, XP055488606, 2 pages.

International Search Report and Written Opinion for PCT Application No. PCT/GB2016/050427, dated May 4, 2016, 15 pages.

Examination Report for European Application No. 16706419.5, dated Jul. 10, 2018, 8 pages.

Long, et al., "Fully Convolutional Networks for Semantic Segmentation", IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 3431-3440.

Wang, et al., "Self-Tuned Deep Super Resolution", 2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops, Jun. 7, 2015, 8 pages.

Zuo, et al., "Content-Adaptive Resolution Enhancement of Compressed Video with Encoder-Generated Side Information", International Conference on Consumer Electronics, Jan. 9, 2008, pp. 1-2.

Office Action for European Application No. 16706418.7, dated Jan. 28, 2019, 6 pages.

Office Action for European Application No. 16715067.1, dated Mar. 6, 2019, 9 pages.

Office Action for European Application No. 16706417.9, dated Mar. 6, 2019, 8 pages.

Office Action for European Application No. 16706424.5, dated Mar. 6, 2019, 8 pages.

Wikipedia, "Universal Approximation Theorem", Jun. 10, 2010, XP055562507, Retrieved from http://web.archive.org/web/20100701194338/https://en.wikipedia.org/wiki/Universal_approximation_theorem, 3 pages.

Atreya, et al., "Novel Lossy Compression Algorithms With Stacked Autoencoders", I available at <http://cs229.stanford.edu/proj2009/AtreyaOshea.pdf>, Dec. 11, 2019; 5 pages.

Bishop, et al., "Super-Resolution Enhancement of Video", International Workshop on Artificial Intelligence and Statistics, no. 9th, Jan. 3, 2003, pp. 1-8.

Kwon, et al., "Adaptive postprocessing algorithm in block-coded images using block classification and MLP", IEEE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, Engineering Sciences Society, vol. E86-A, No. 4, Apr. 1, 2003, pp. 961-967.

Nielsen, "Neural Networks and Deep Learning, Chapter 4: A Visual Proof That Neural Nets Can Compute Any Function," Feb. 6, 2015, XP055562061, 32 pages.

Park, et al., "Super-Resolution Image Reconstruction: A Technical Overview", IEEE Signal Processing Magazine, IEEE Service Center, vol. 20, No. 3, May 1, 2003, pp. 21-36.

Yang, et al., "Single-Image Super-Resolution Reconstruction via Learned Geometric Dictionaries and Clustered Sparse Coding", IEEE Transactions on Image Processing, vol. 21, No. 9, Sep. 1, 2012, pp. 4016-4028.

Examination Report for European Application No. 16706419.5, dated Mar. 6, 2019, 6 pages.

(56)

References Cited

OTHER PUBLICATIONS

Bevilacqua, et al., "Video Super-Resolution Via Sparse Combinations of Key-Frame Patches in a Compression Context", 30th Picture Coding Symposium (PCS), Dec. 8, 2013, 5 pages.

Cheng, et al., "Fast Video Super-Resolution Using Artificial Neural Networks", 8th IEEE, IET International Symposium on Communication Systems, Networks and Digital Signal Processing, Jul. 18, 2012, 4 pages.

Dong, "Image Super-Resolution Using Deep Convolutional Networks", IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 38, No. 2, Feb. 1, 2016, 14 pages.

Gendy, et al., "A Fast Algorithm for Image Restoration Using a Recurrent Neural Network with Bound-Constrained Quadratic Optimization", The Seventh Australian and New Zealand Intelligent Information Systems Conference, Nov. 18, 2001, pp. 111-116.

Liang, et al., "Incorporating Image Degeneration Modeling With Multitask Learning for Image Super-Resolution", IEEE International Conference on Image Processing, Sep. 27, 2015, pp. 2110-2114.

Liu, et al., "An Adaptive Image Denoising Model Based on Tikhonov and TV Regularizations", Advances in Multimedia, vol. 2014, Jan. 1, 2014, 11 pages.

Muneyasu, et al., "Image Restoration Using Layered Neural Networks and Hopfield Networks", Proceedings of the International Conference on Image Processing, vol. 2, Oct. 23, 1995, pp. 33-36.

Qu, et al., "Single Image Super-Resolution via Convolutional Neural Network and Total Variation Regularization", MultiMedia

Modeling. Lecture Notes in Computer Science, vol. 951, Jan. 1, 2016, pp. 28-38.

Reed, et al., "Generative Adversarial Text to Image Synthesis", Proceedings of the 33rd International Conference on Machine Learning, JMLR: W&CP vol. 48, May 2016, 9 pages.

Rudin, et al., "Non-Linear Total Variation Based Noise Removal Algorithms", Physica D, vol. 60, No. 1-4, Nov. 1, 1992, pp. 259-268.

Sheikh, et al., "An Information Fidelity Criterion for Image Quality Assessment Using Natural Scene Statistics", IEEE Transactions on Image Processing, vol. 14, No. 12, Dec. 1, 2005, pp. 2117-2128.

Shi, et al., "Is the Deconvolution Layer the Same As a Convolutional Layer?", retrieved on Mar. 11, 2019 from <https://arxiv.org/ftp/arxiv/papers/1609/1609.07009.pdf>, Sep. 22, 2016, 7 pages.

Sumali, et al., "Single Image Super Resolution by No-Reference Image Quality Index Optimization in PCA Subspace", IEEE 12th International Colloquium on Signal Processing and Its Applications, Mar. 4, 2016, pp. 183-187.

Wang et al., "Deep Networks for Image Super-Resolution with Sparse Prior", International Conference on Computer Vision (ICCV), Dec. 1, 2015, 9 pages.

Xie, et al., "Image Denoising and Inpainting With Deep Neural Networks", Advances in neural information processing systems, 2012, 9 pages.

Yu, et al., "VLSI Friendly Fast CU/PU Mode Decision for HEVC Intra Encoding: Leveraging Convolution Neural Network", 2015 IEEE International Conference on Image Processing, Sep. 27, 2015, pp. 1285-1289.

* cited by examiner




Fig. 1




Fig. 2a





Fig. 2b

Fig. 3






Fig. 4

Fig. 5





Fig. 6



700

Fig. 7




Fig. 8

800





Fig. 9



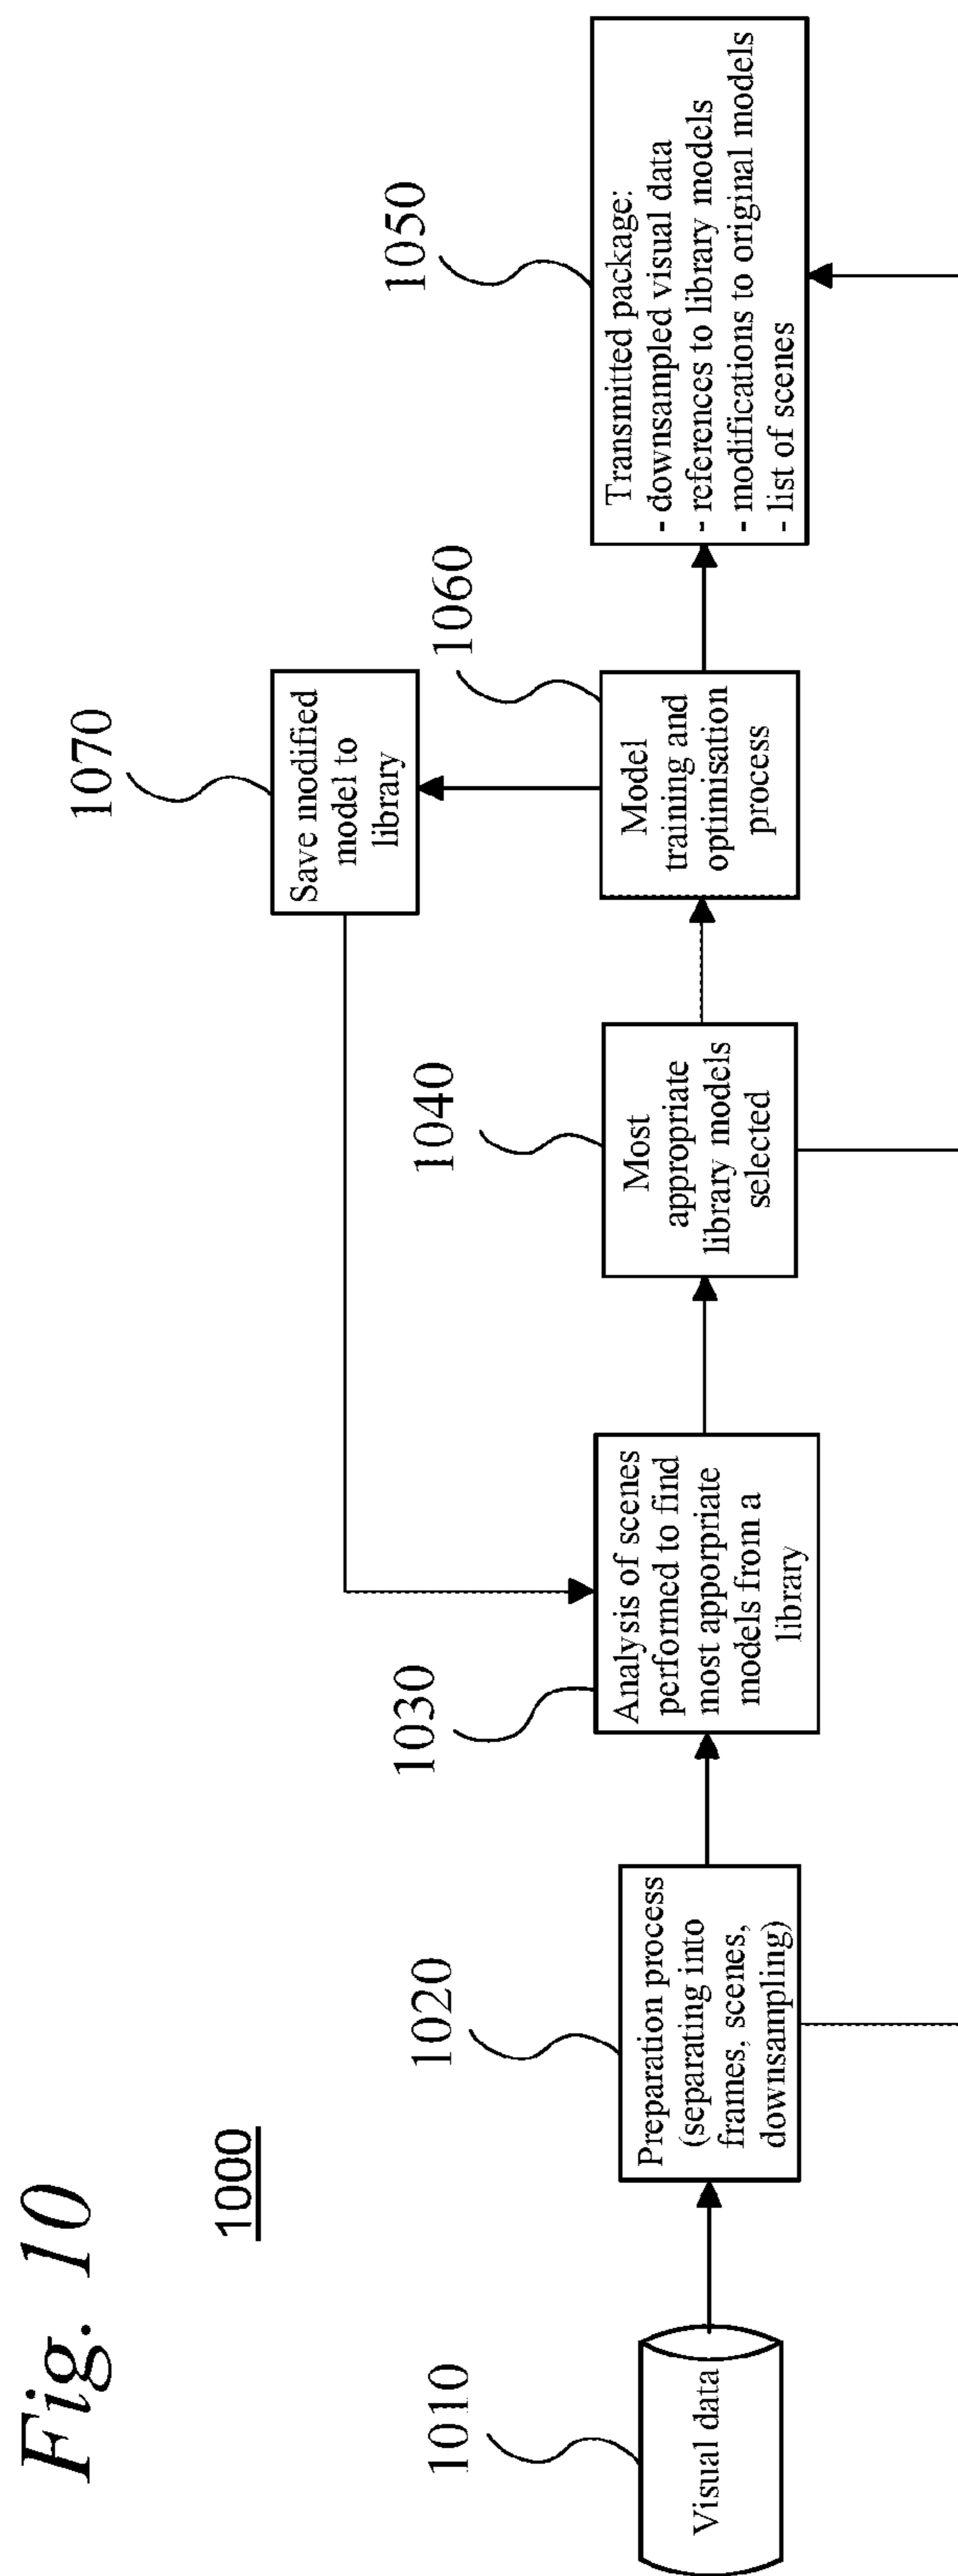


Fig. 11

1100

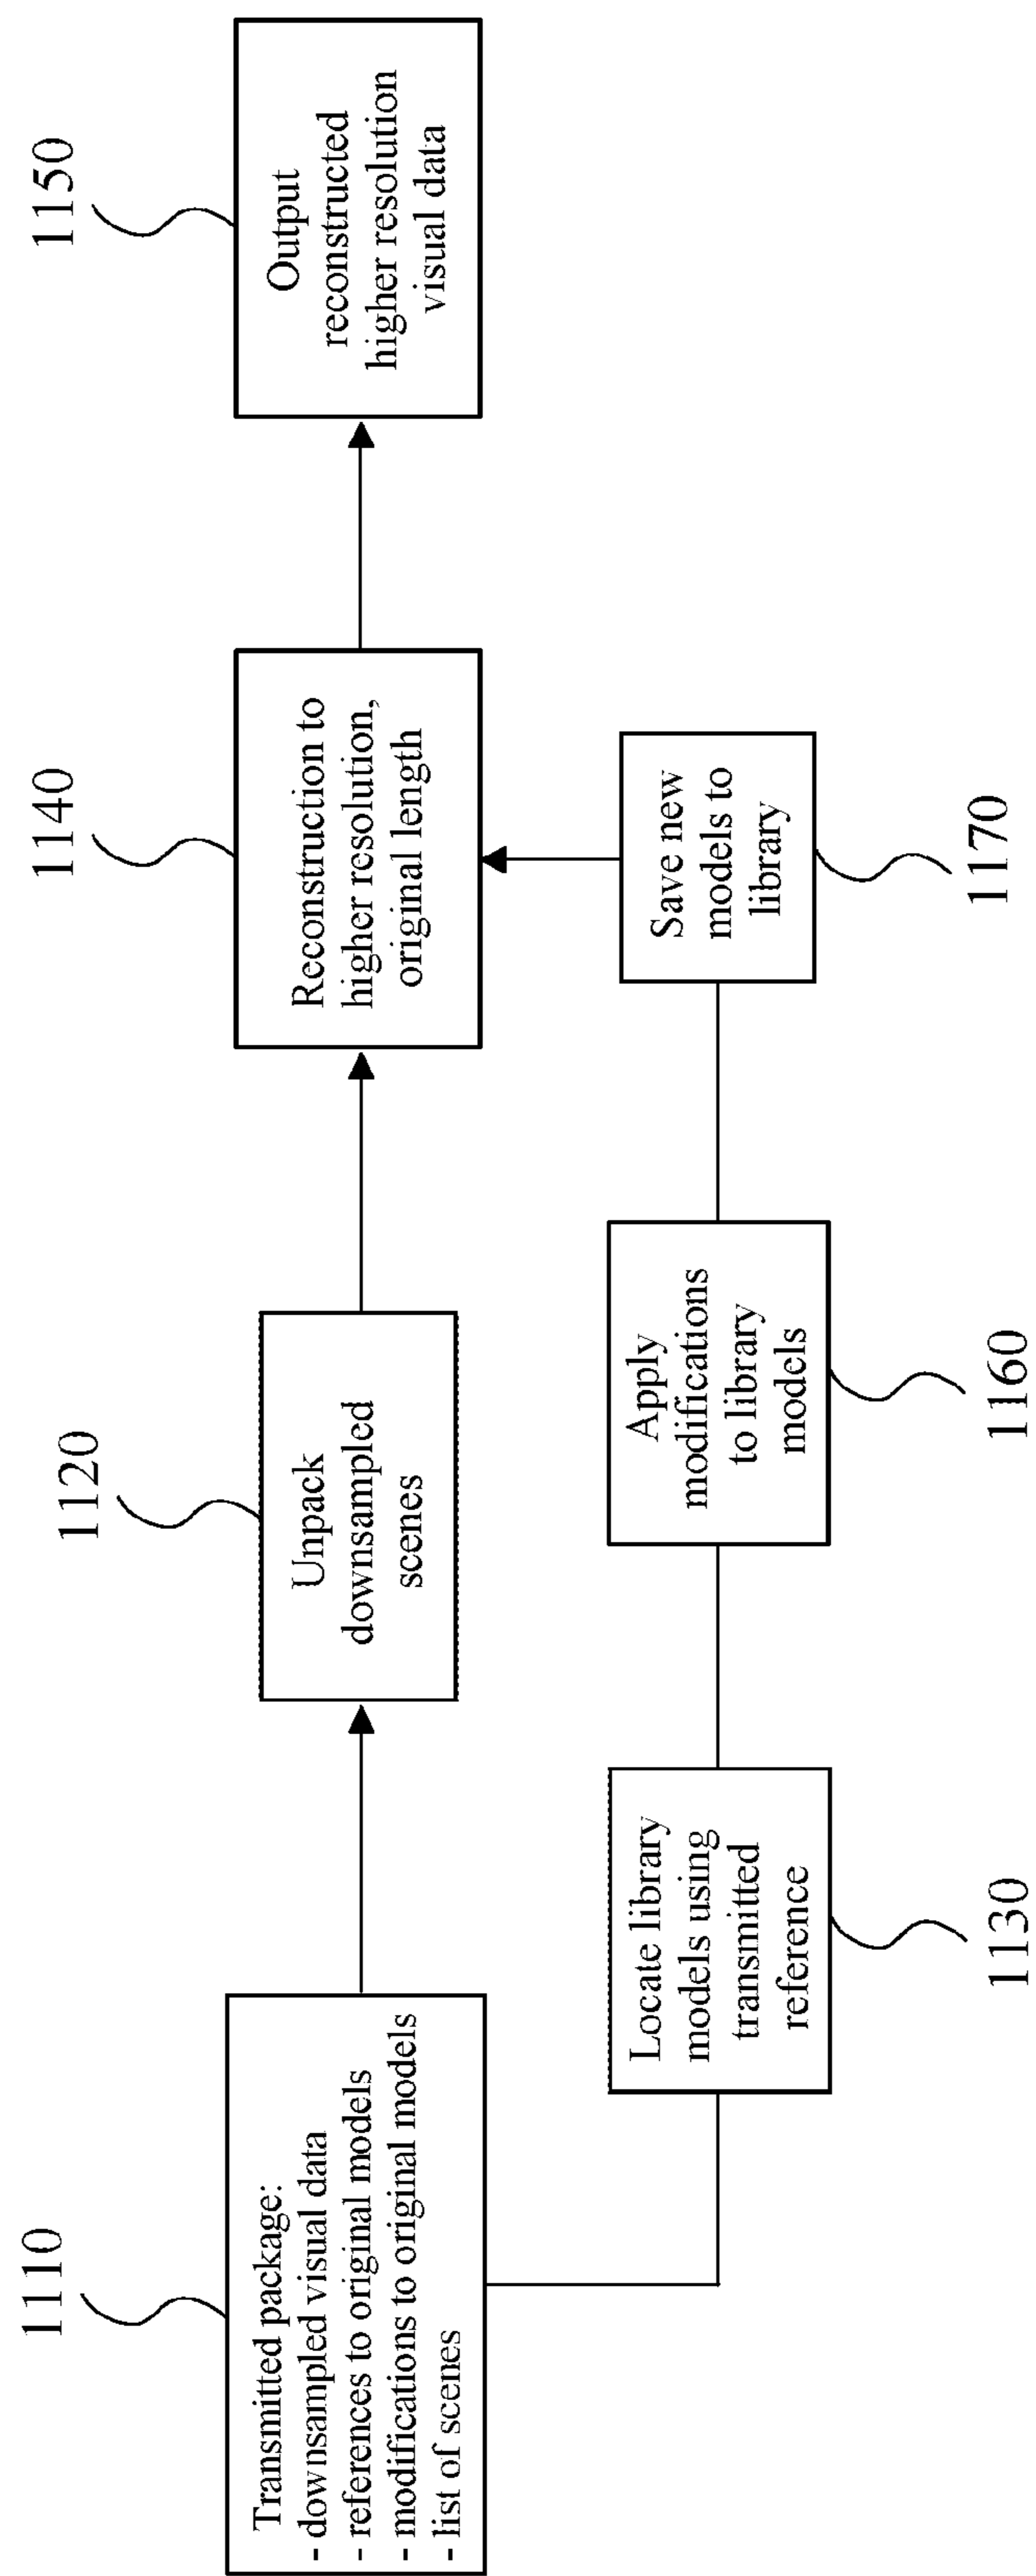


Fig. 12

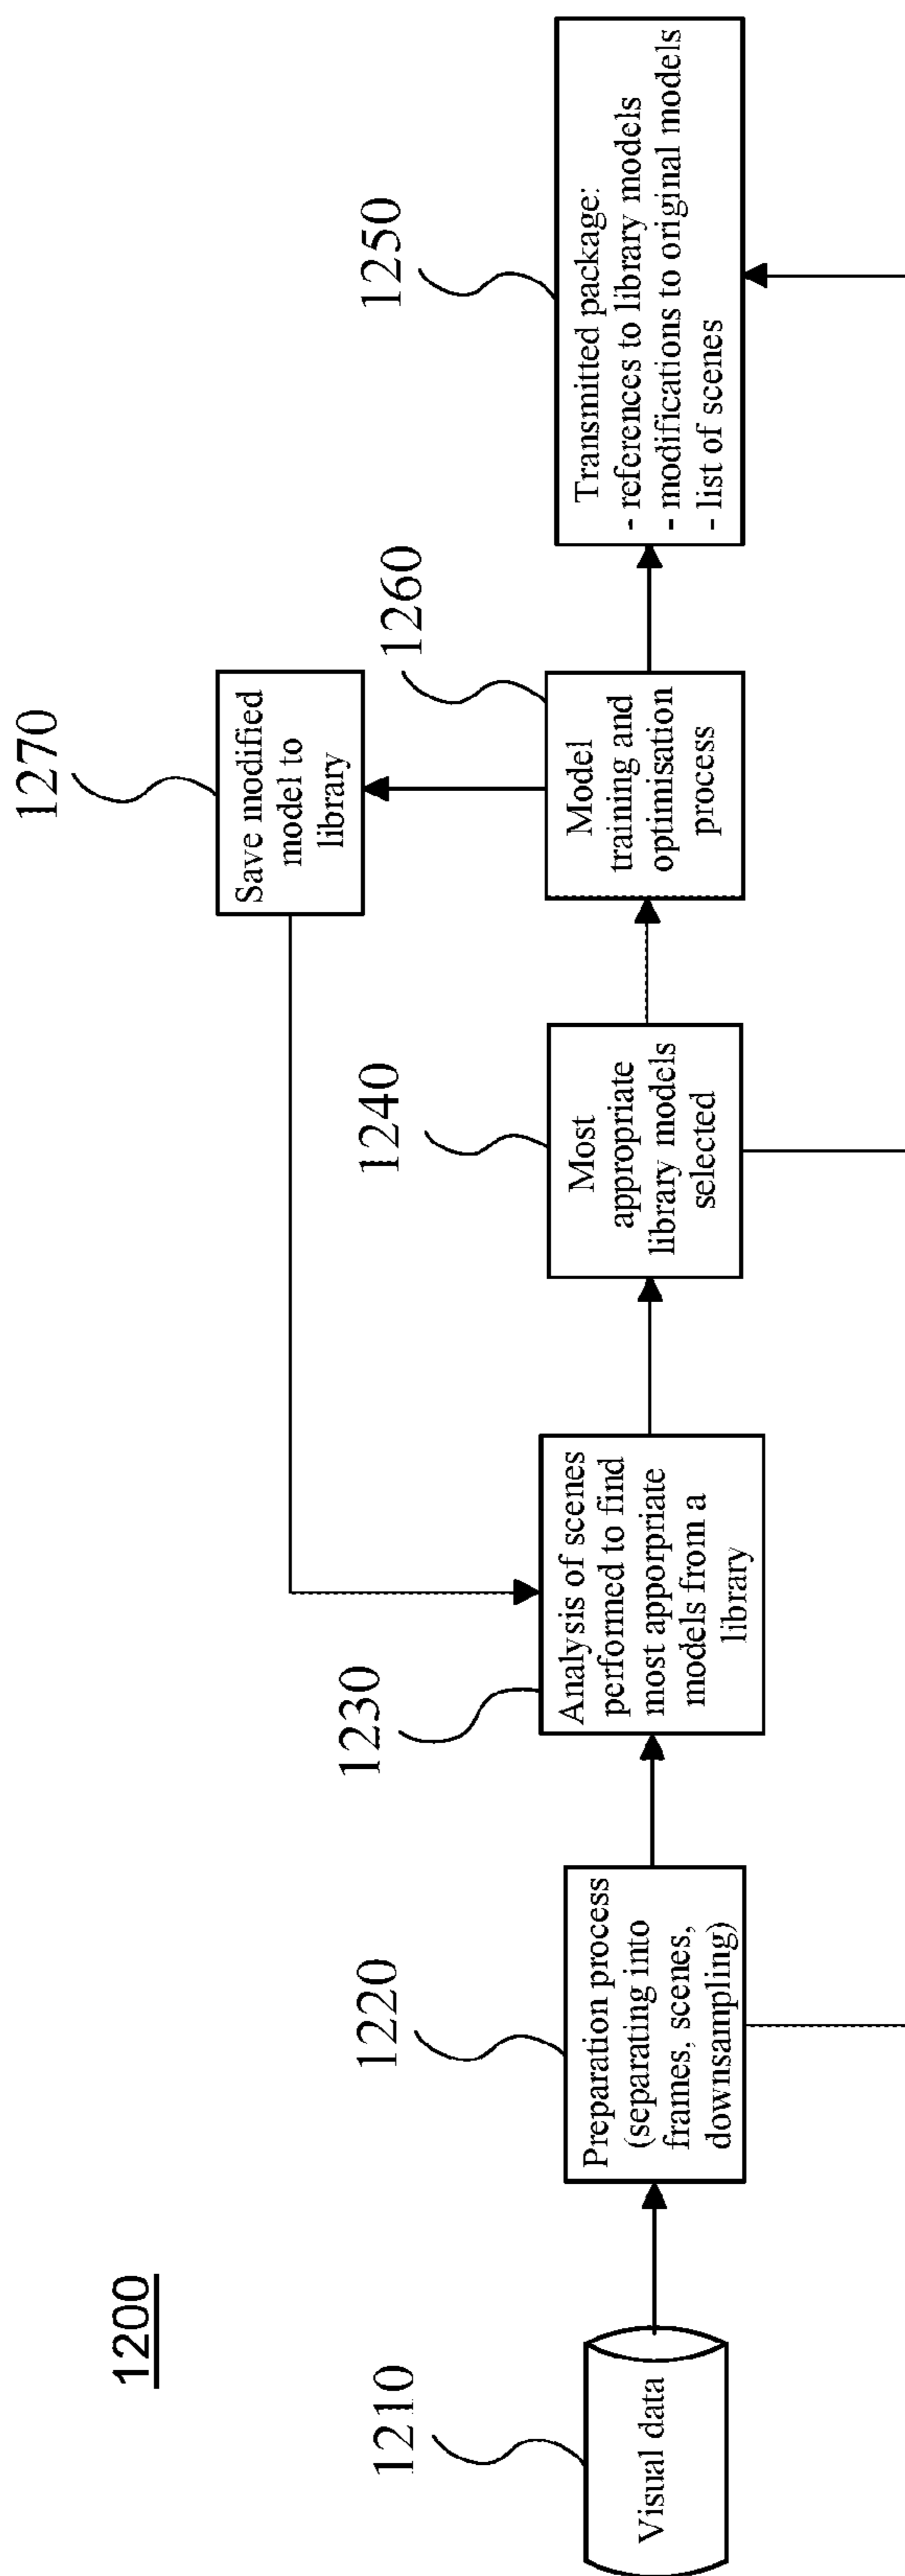


Fig. 13

1300




Fig. 14

1400




Fig. 15




Fig. 16

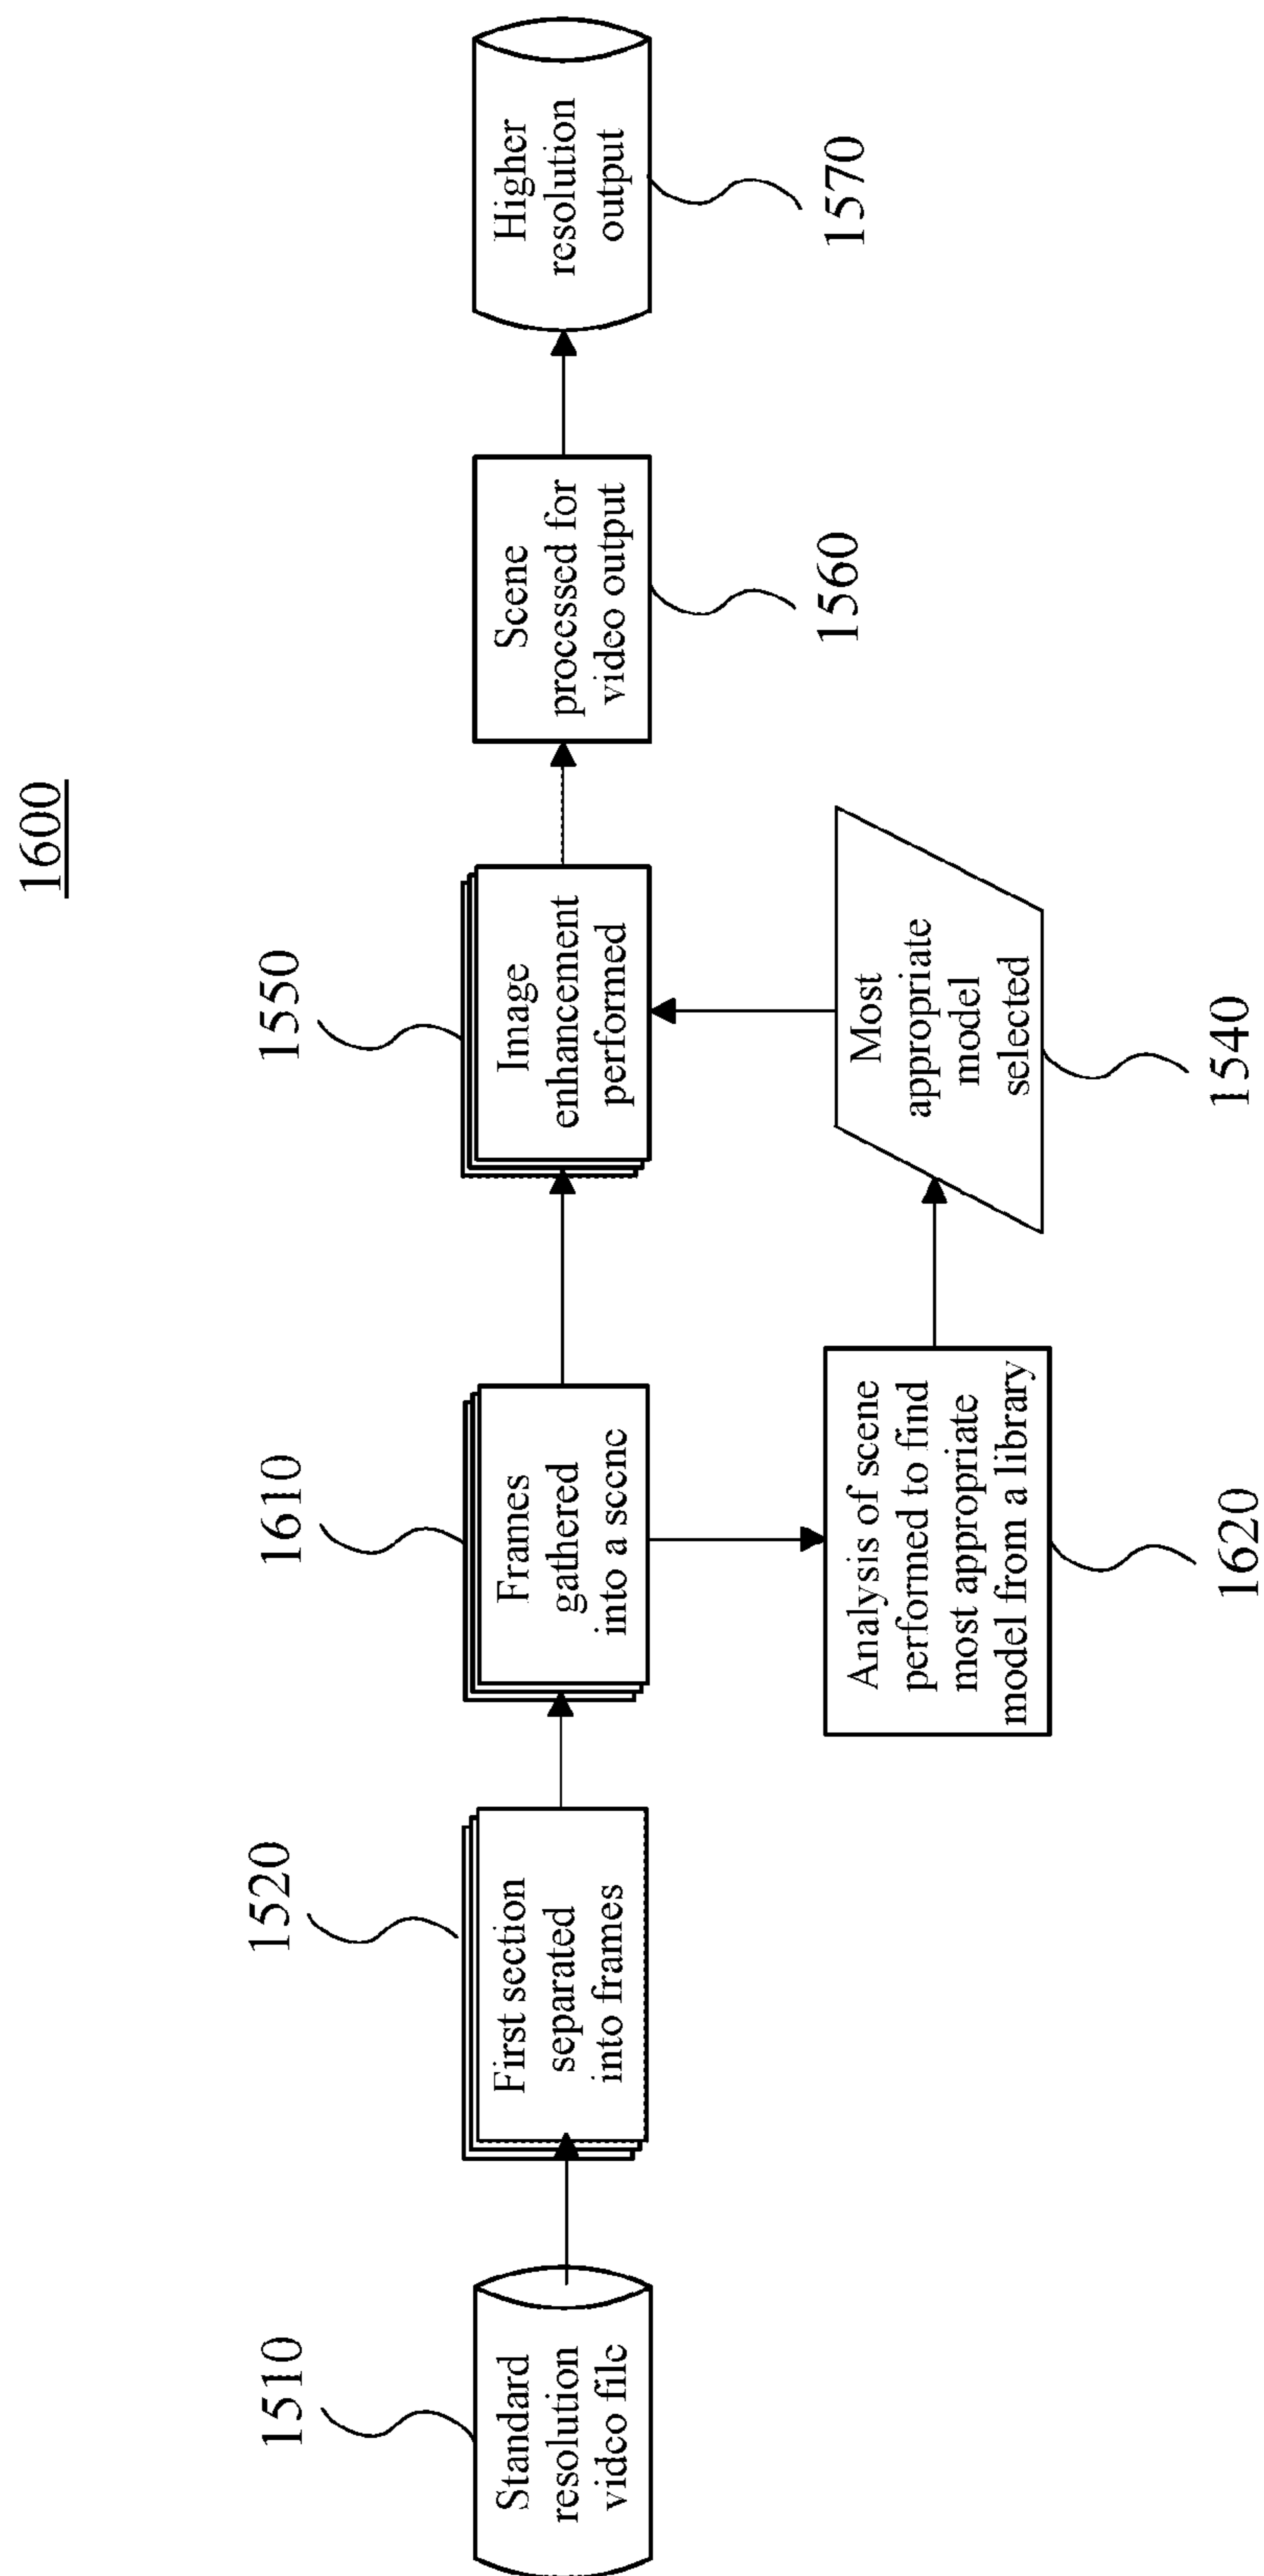


Fig. 17




Fig. 18

1800

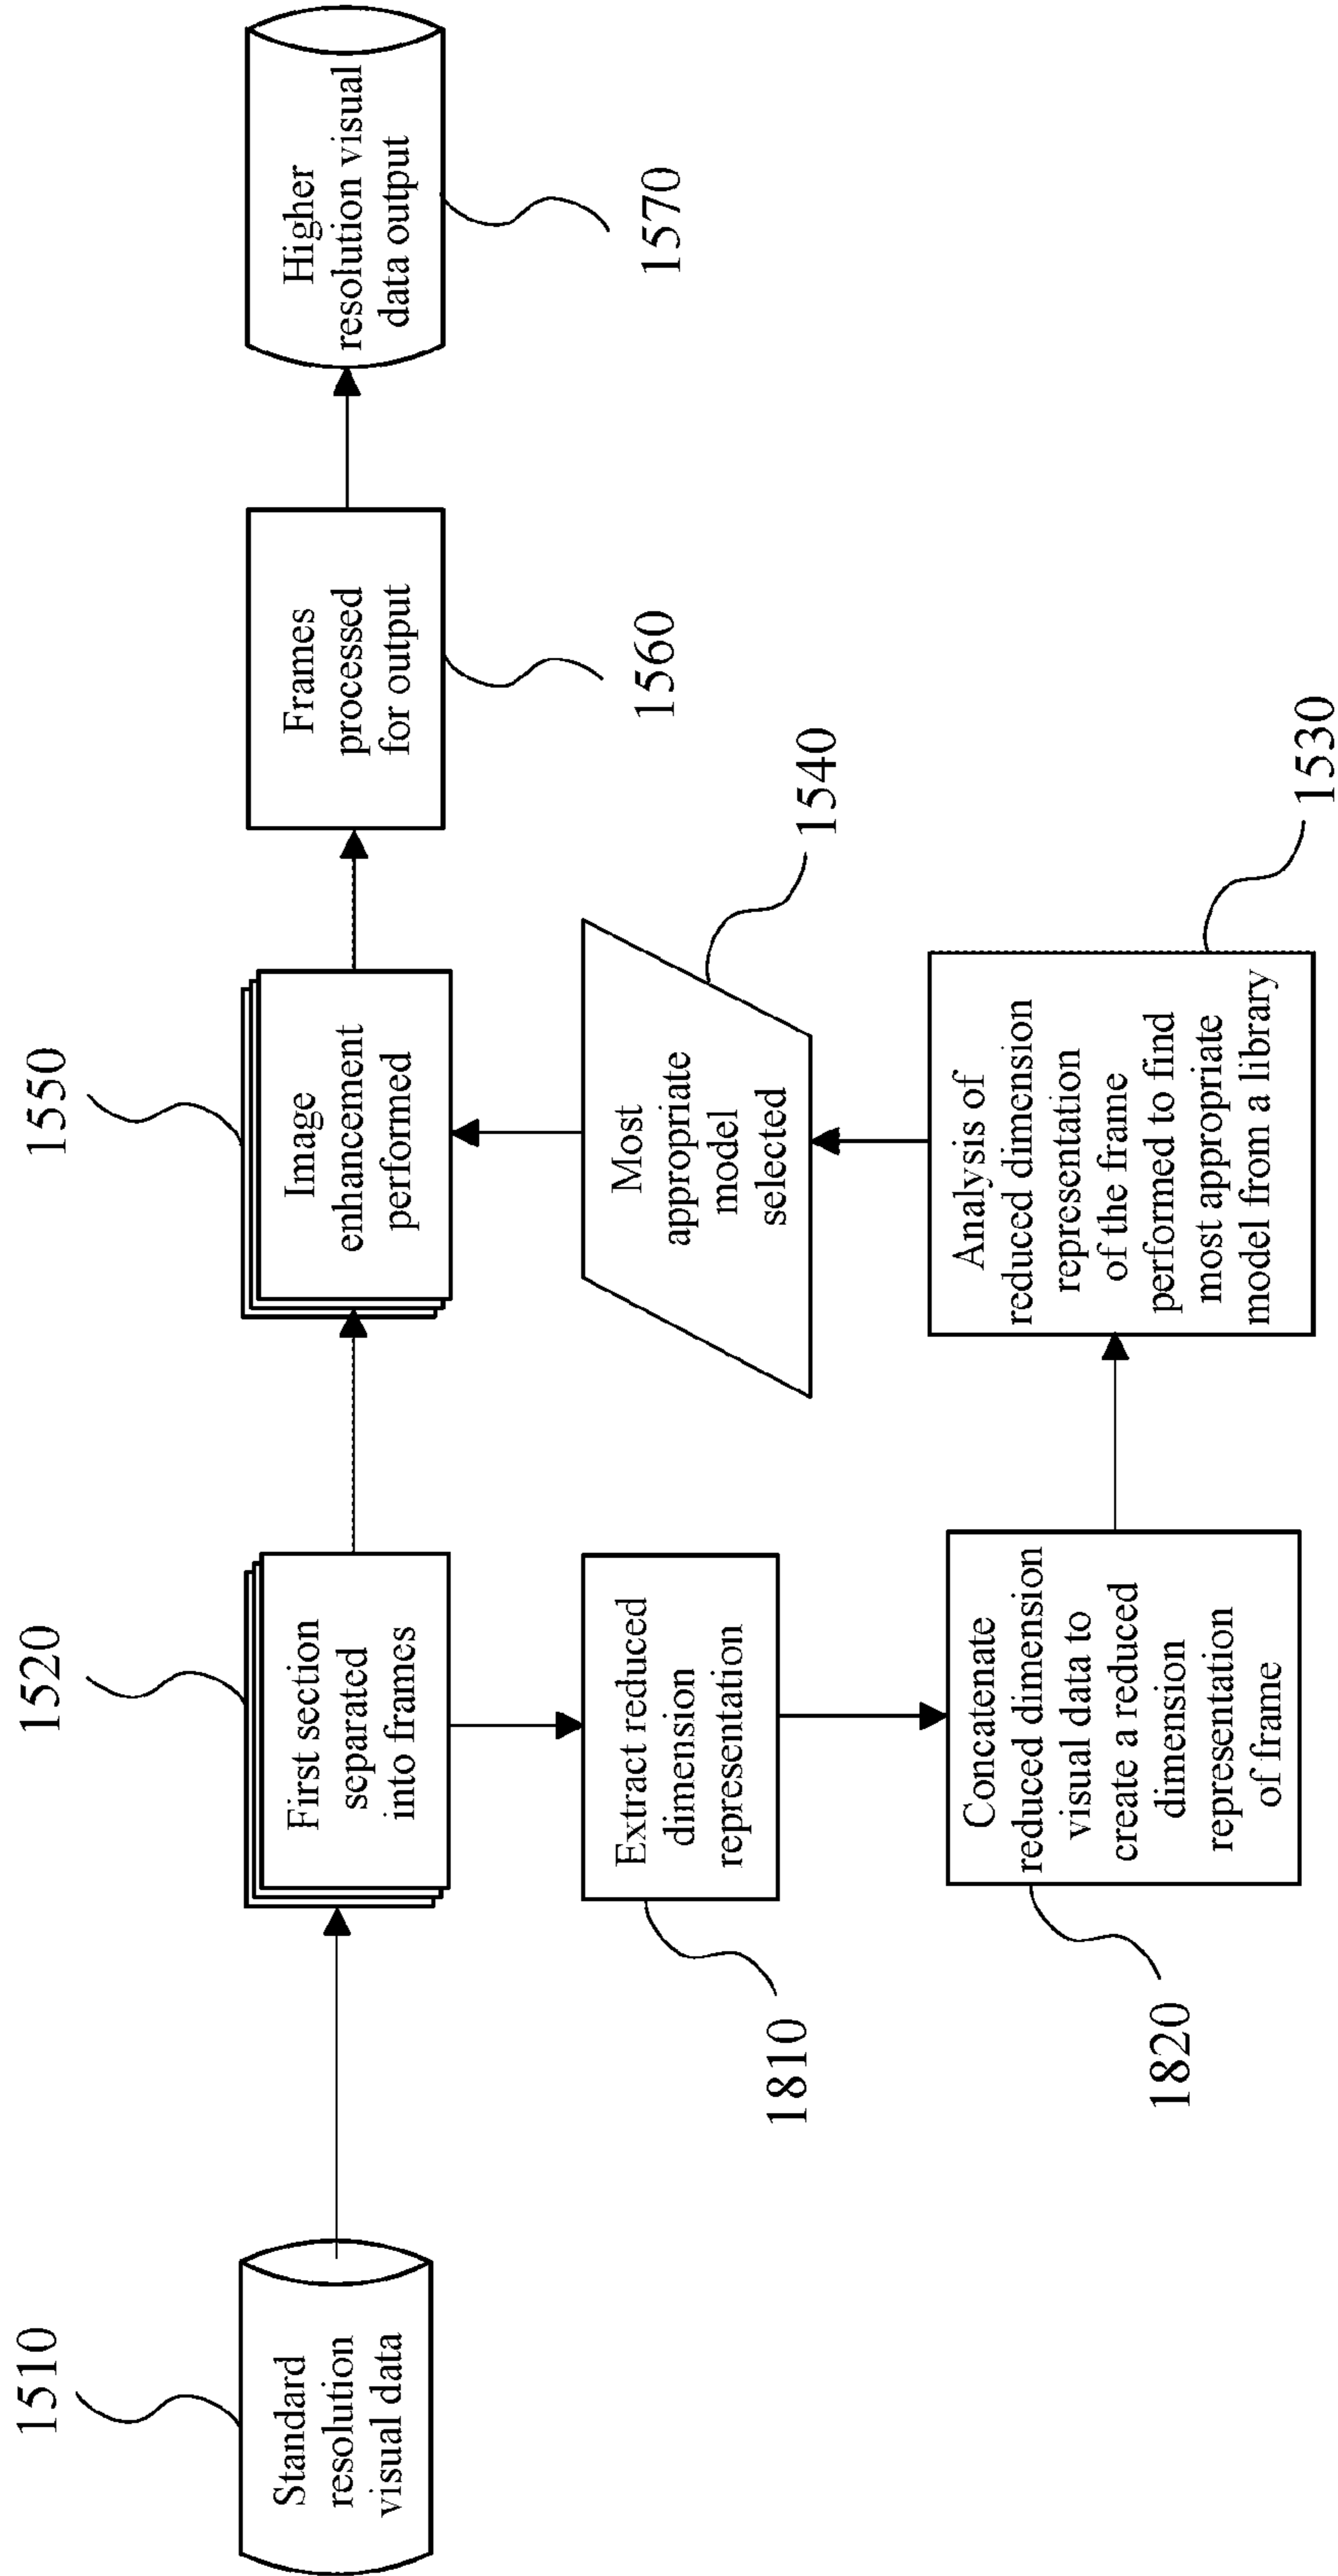


Fig. 19

1900

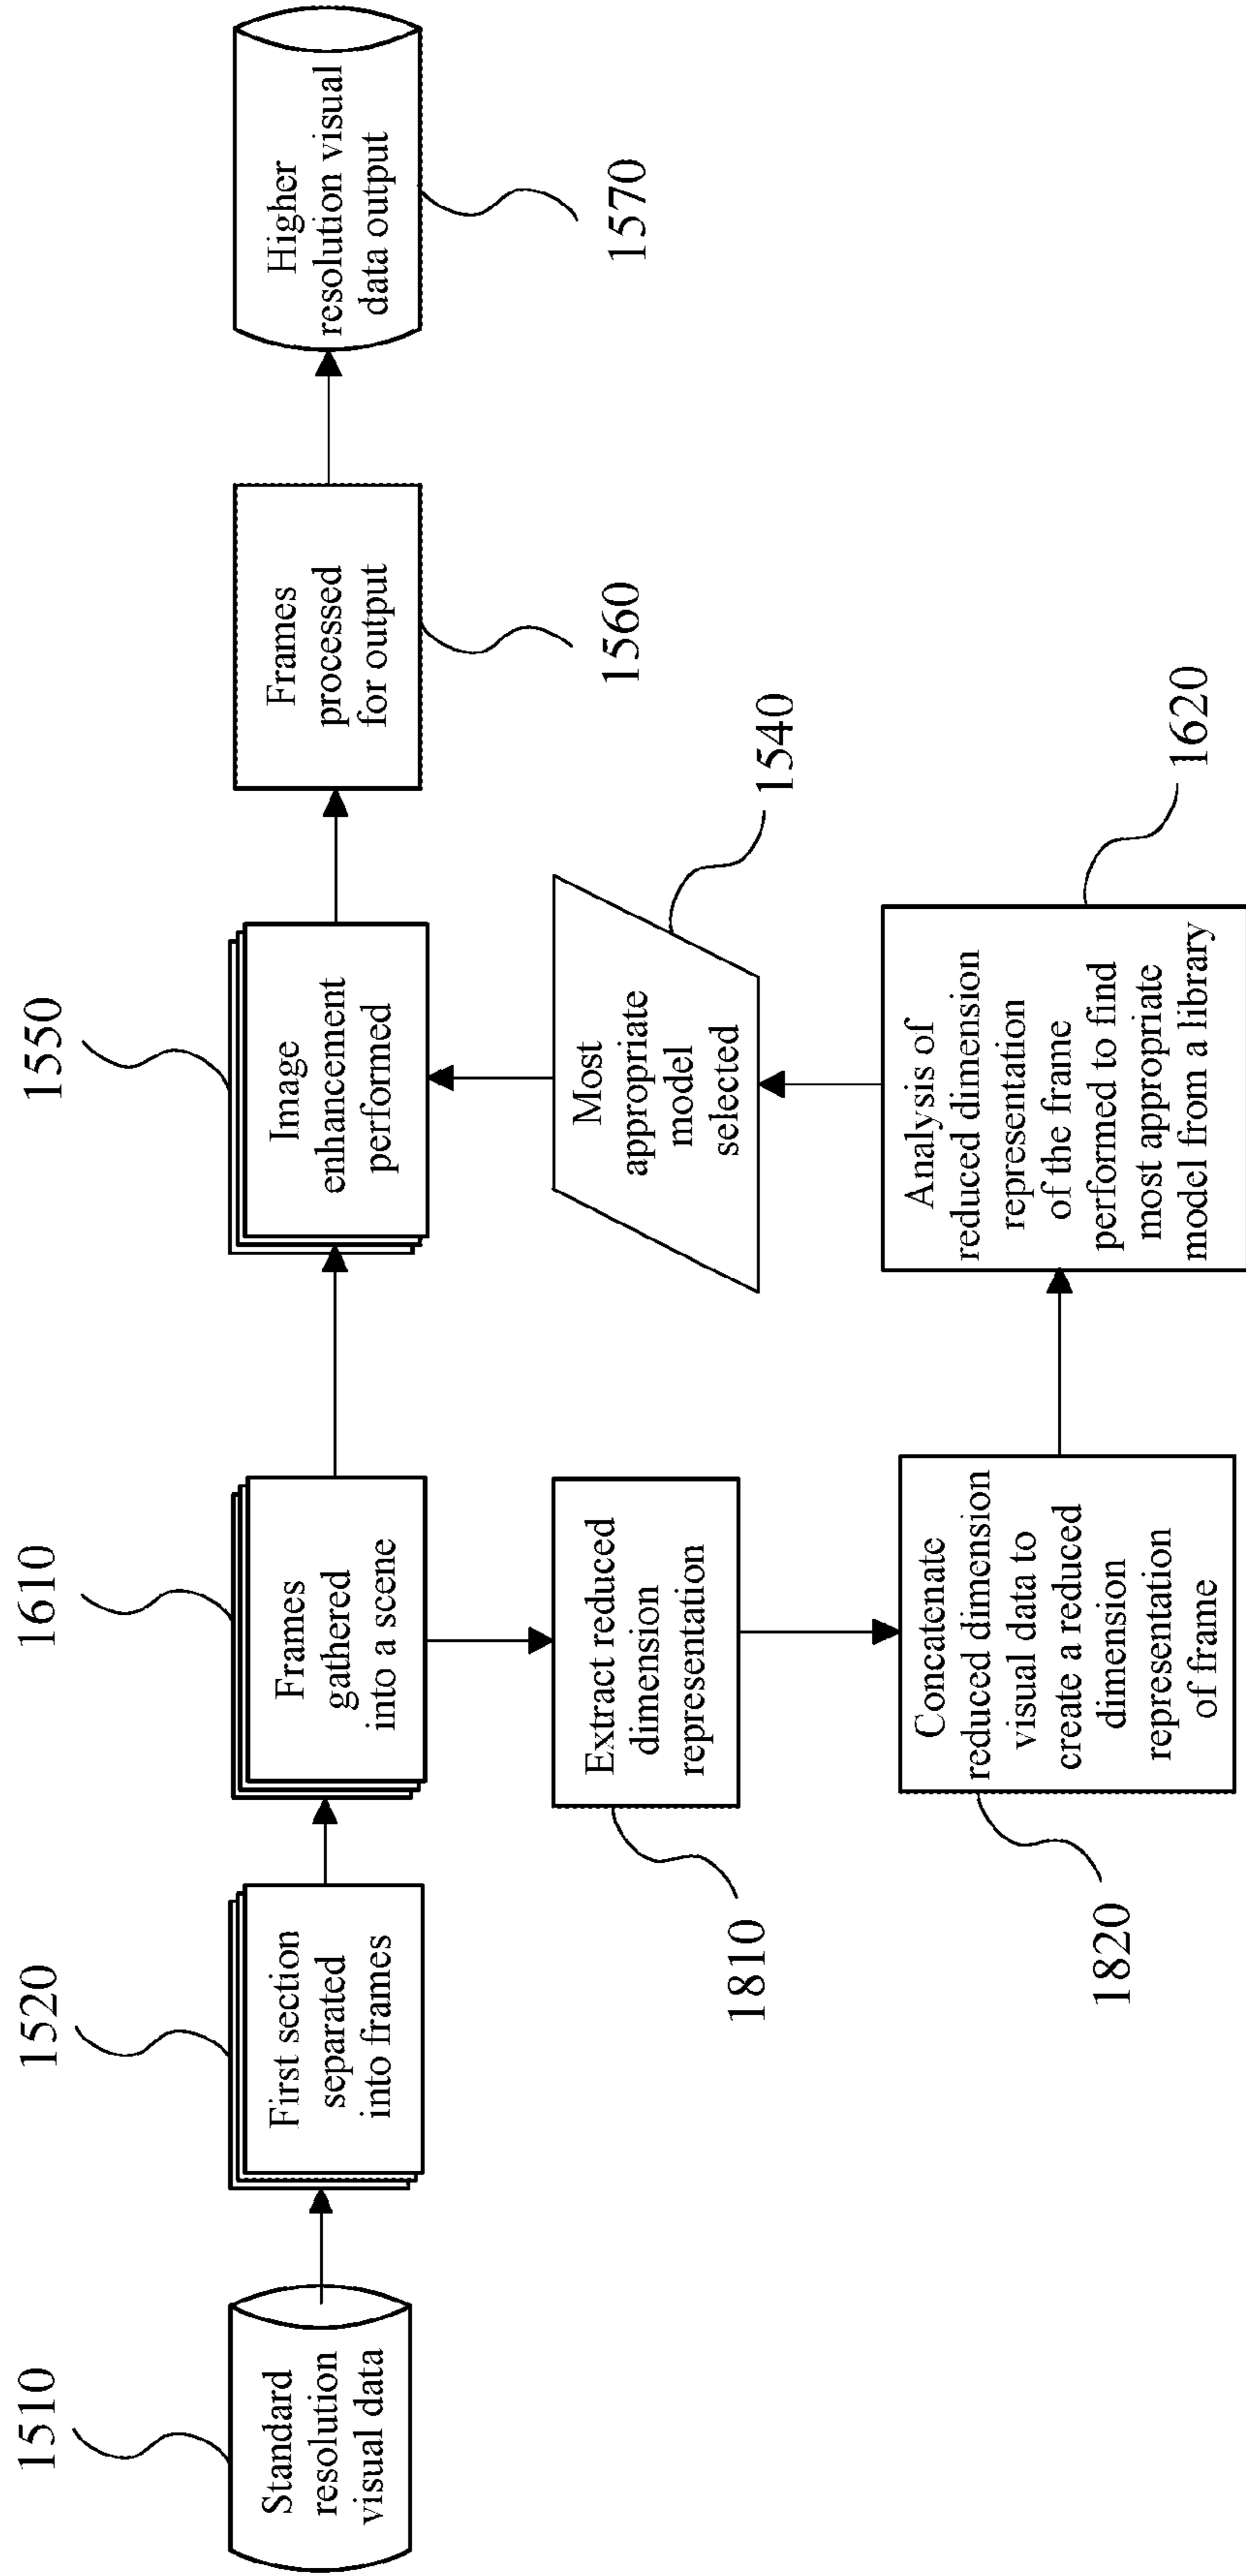


Fig. 20

2000

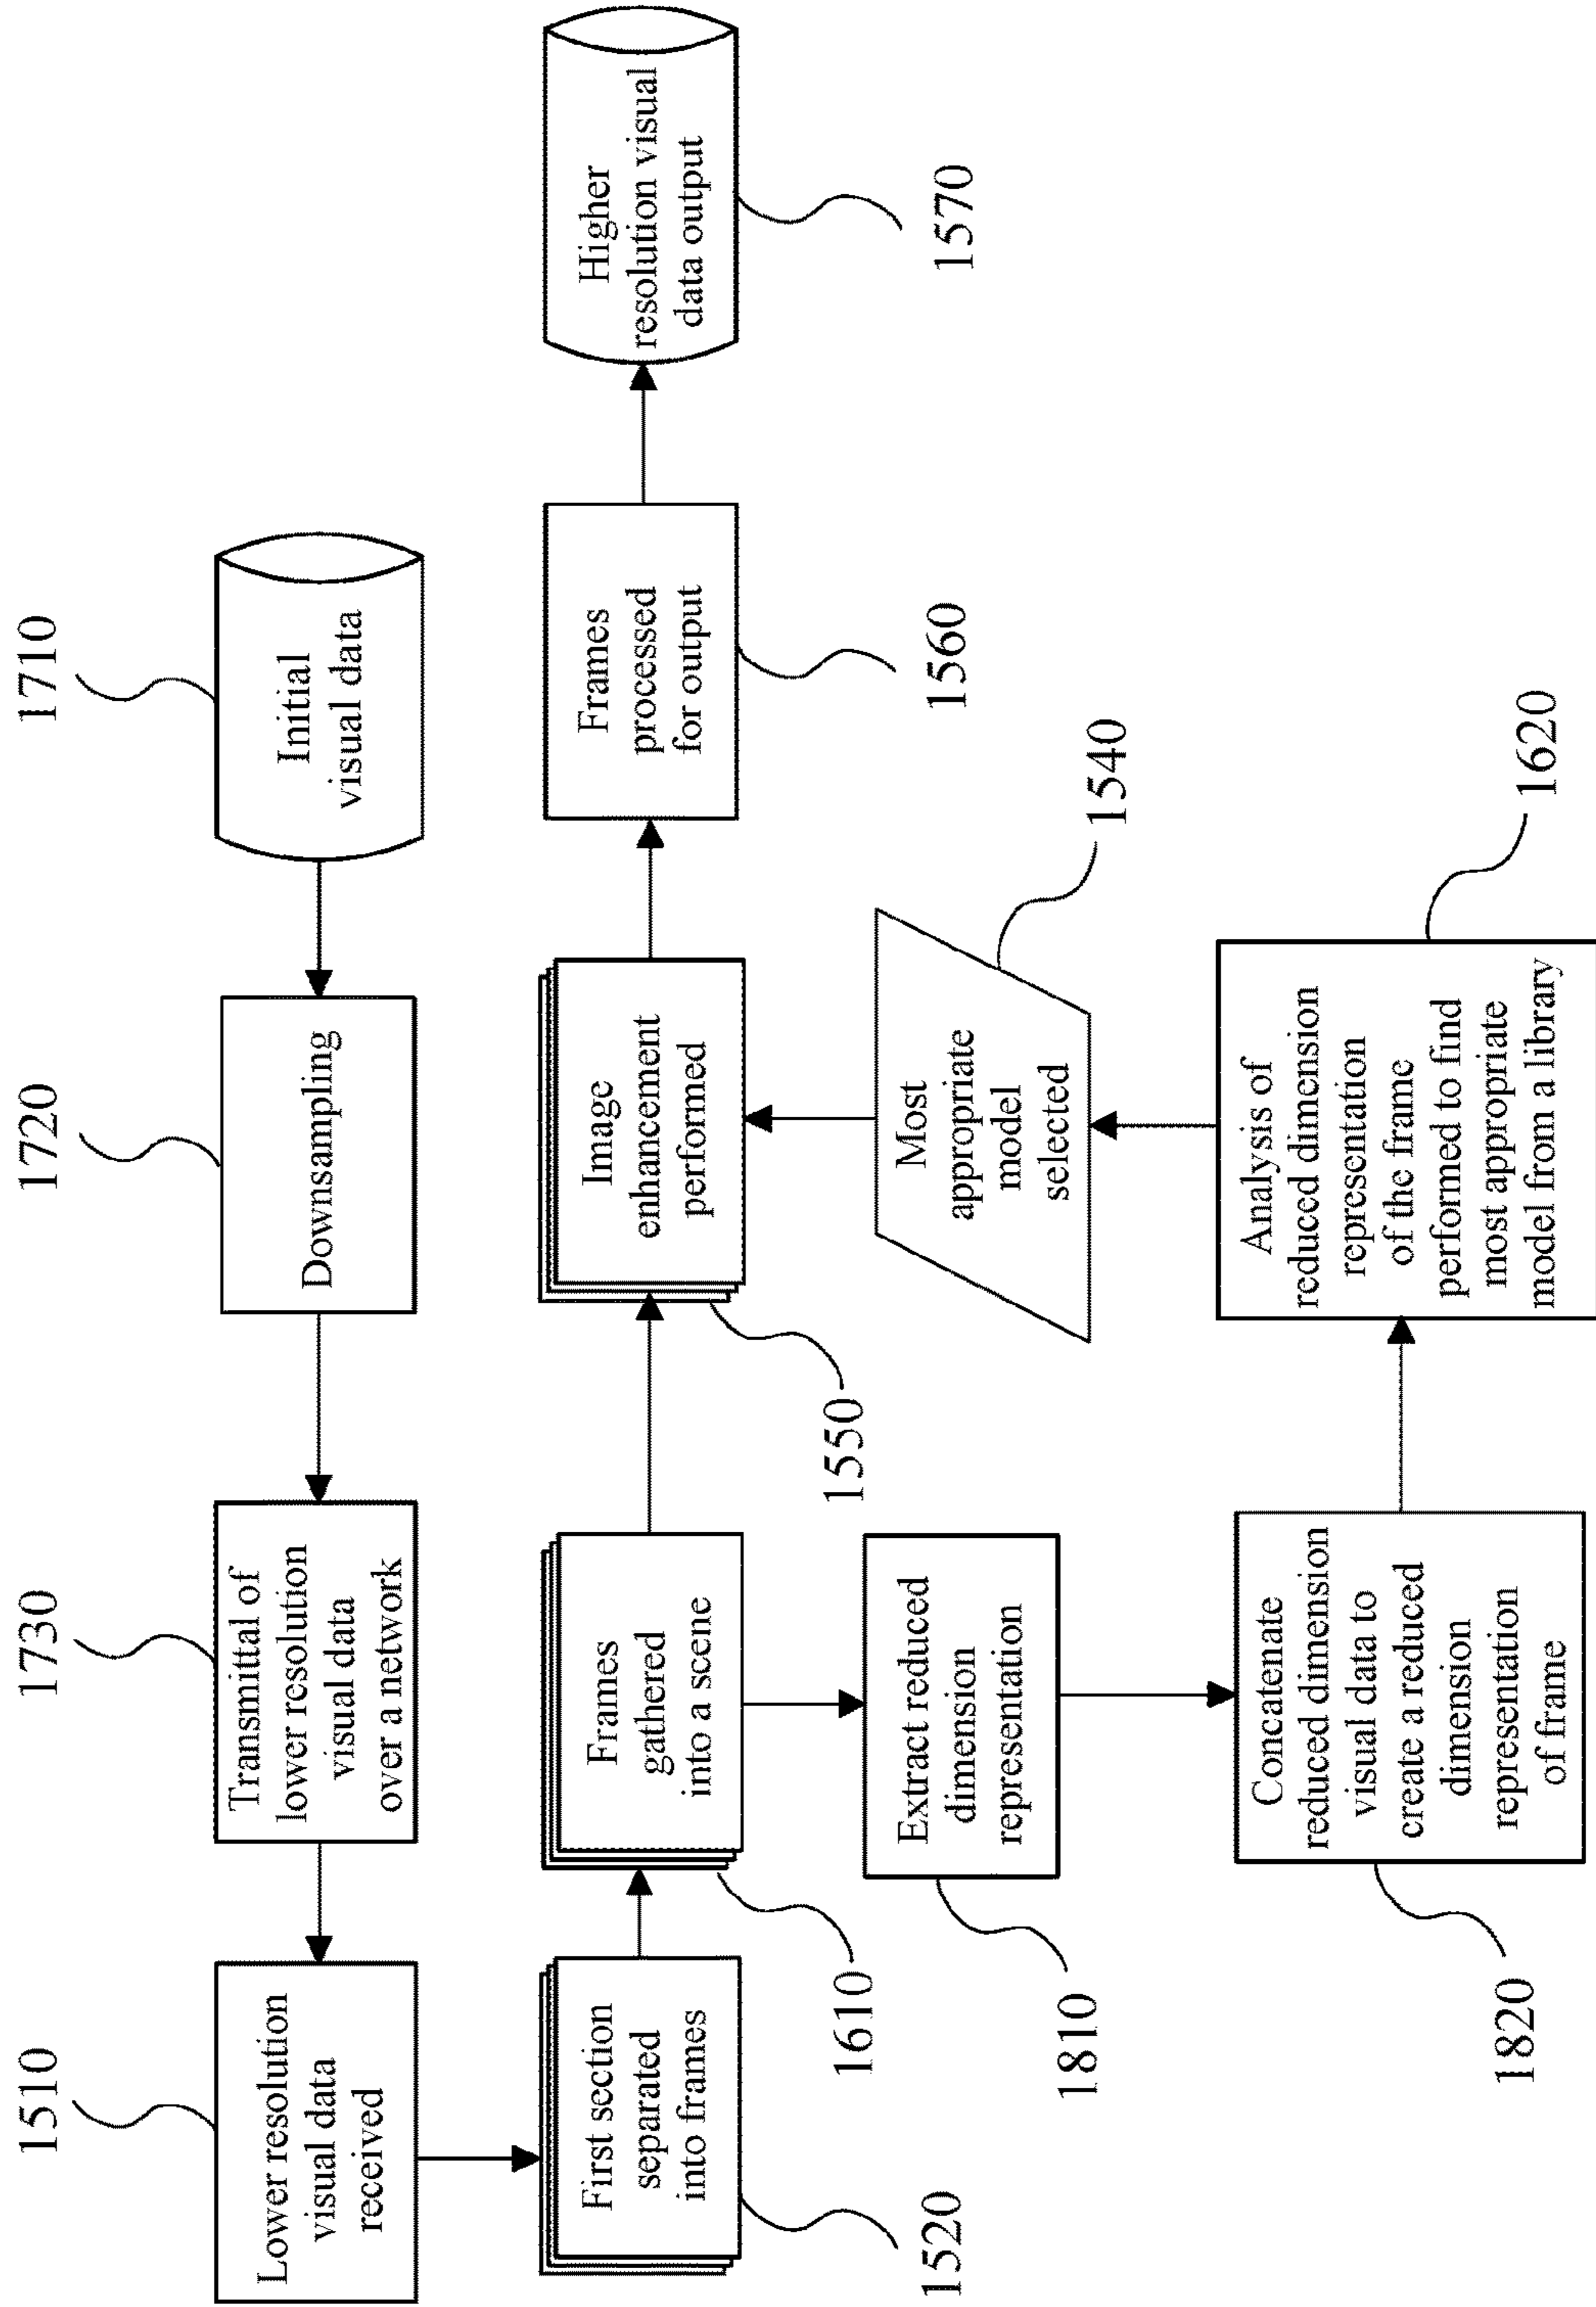


Fig. 21

2100




Fig. 22

2200




Fig. 23

2300

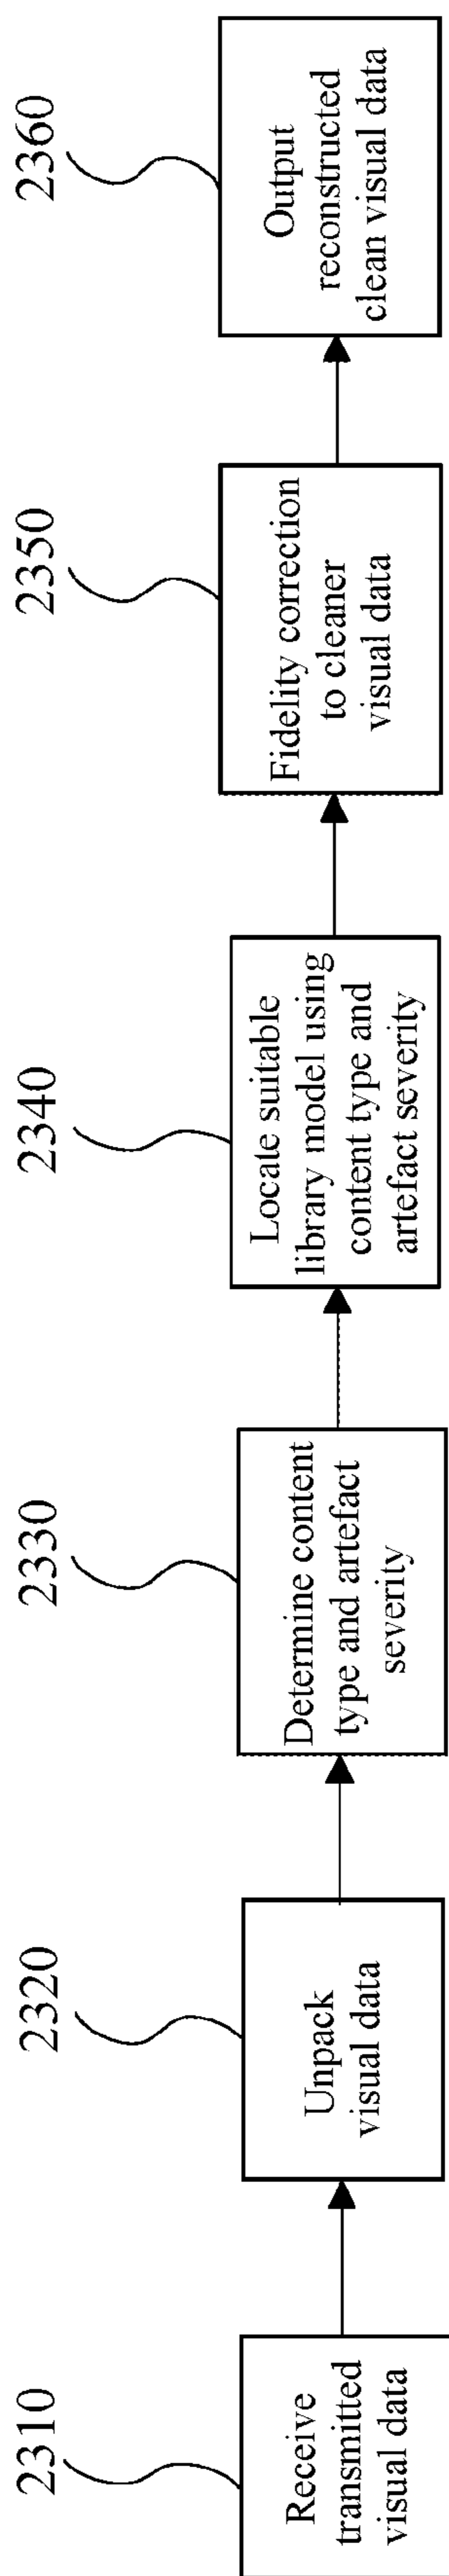


Fig. 24

2400

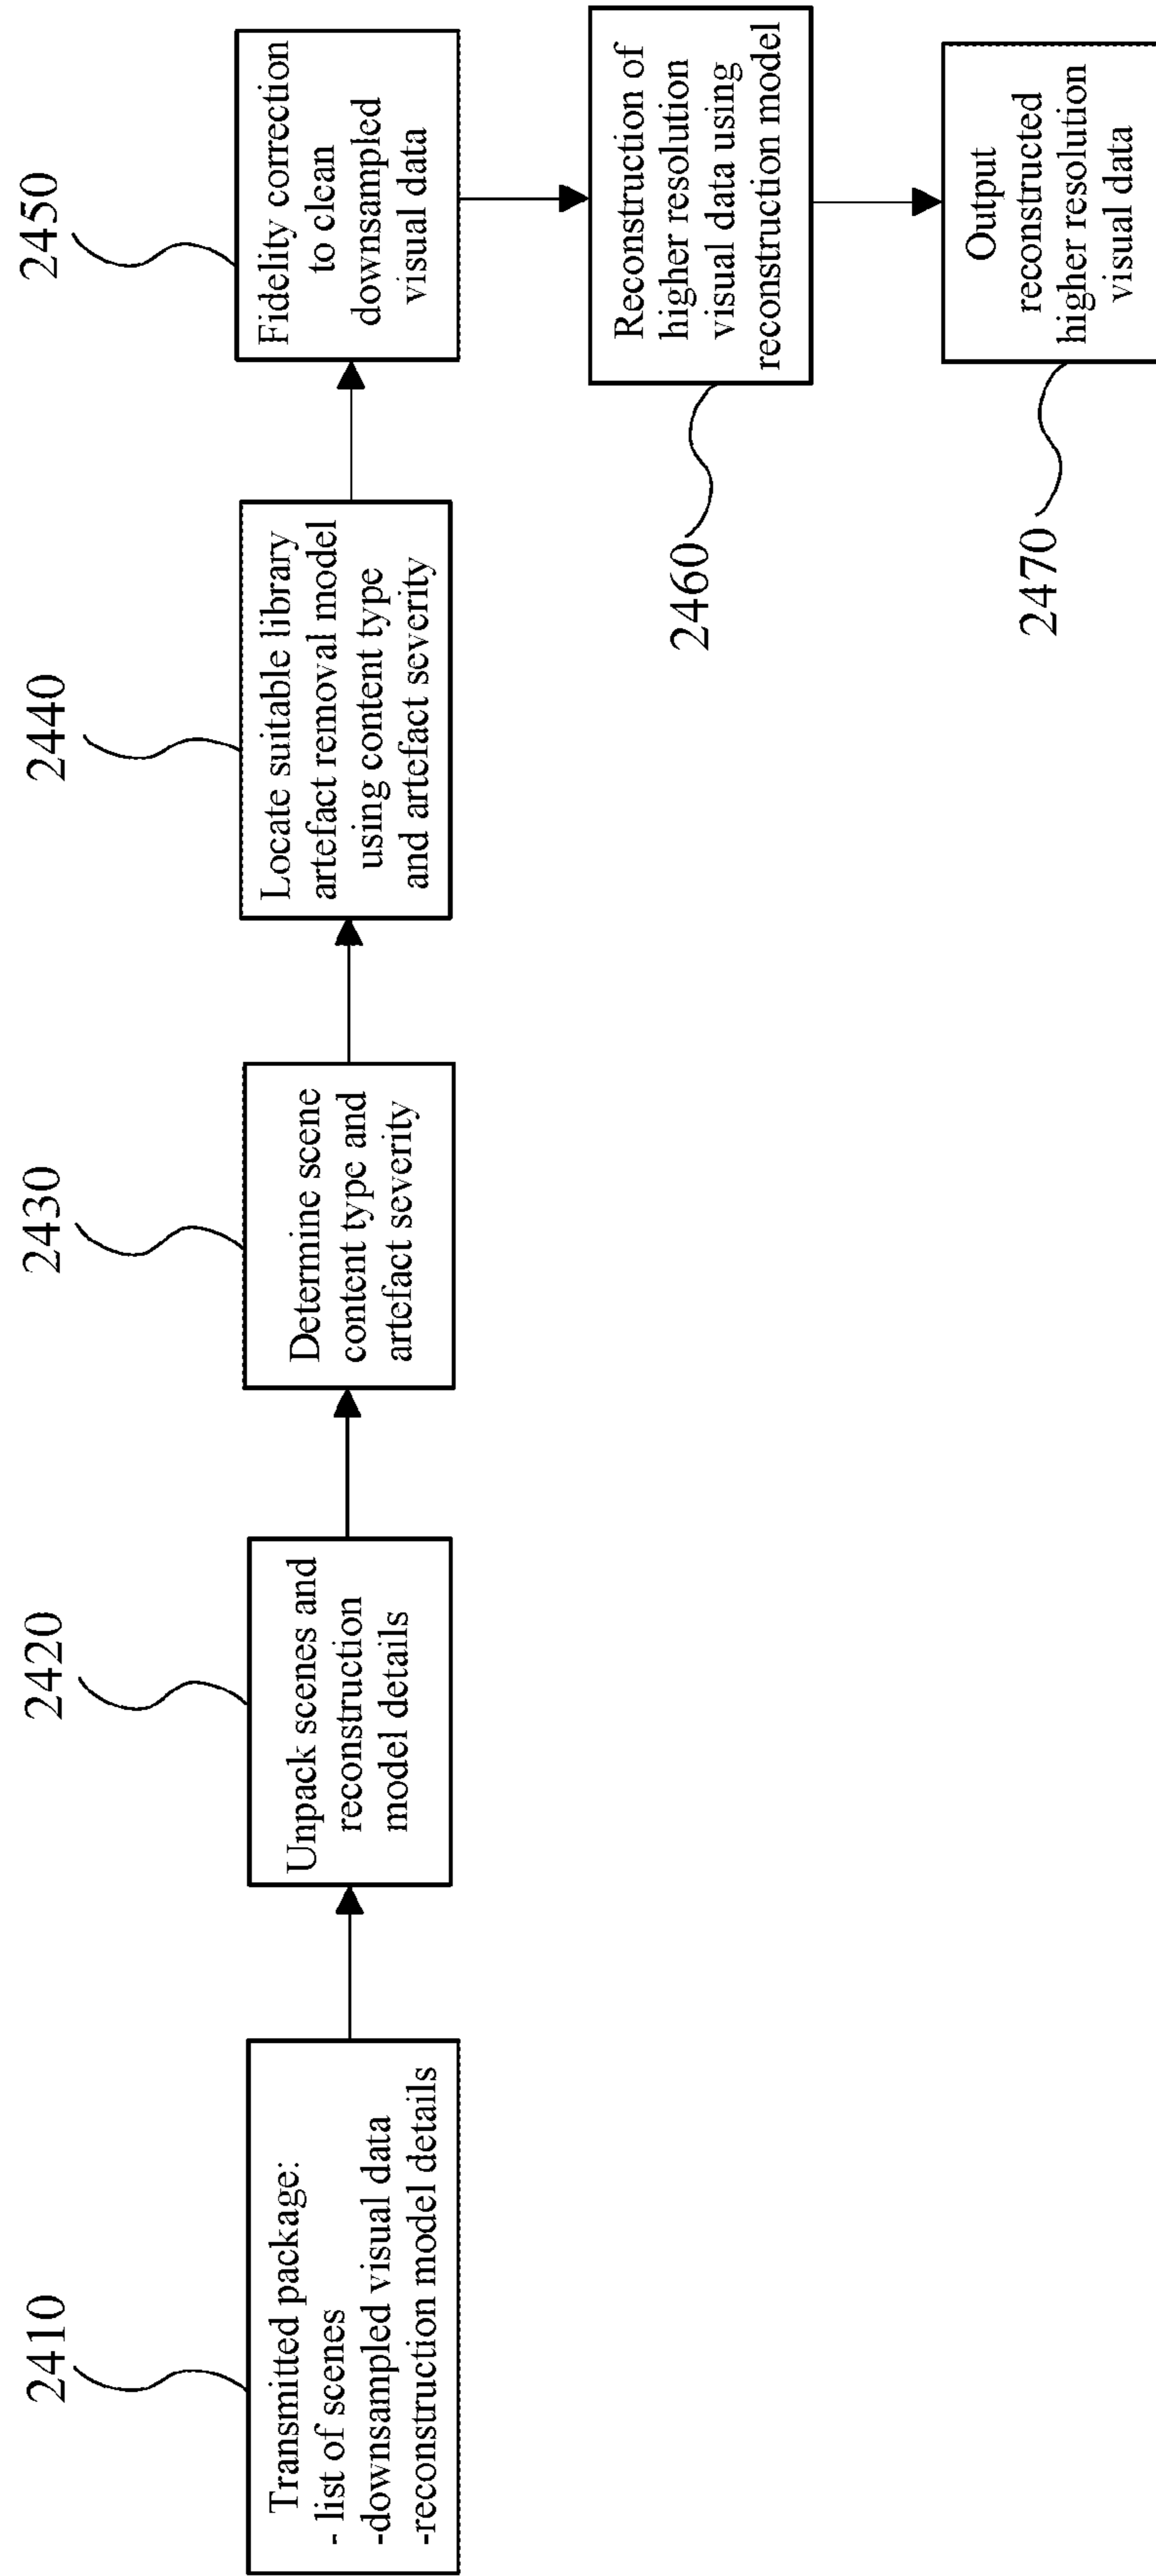



Fig. 25

2500



2600

Fig. 26

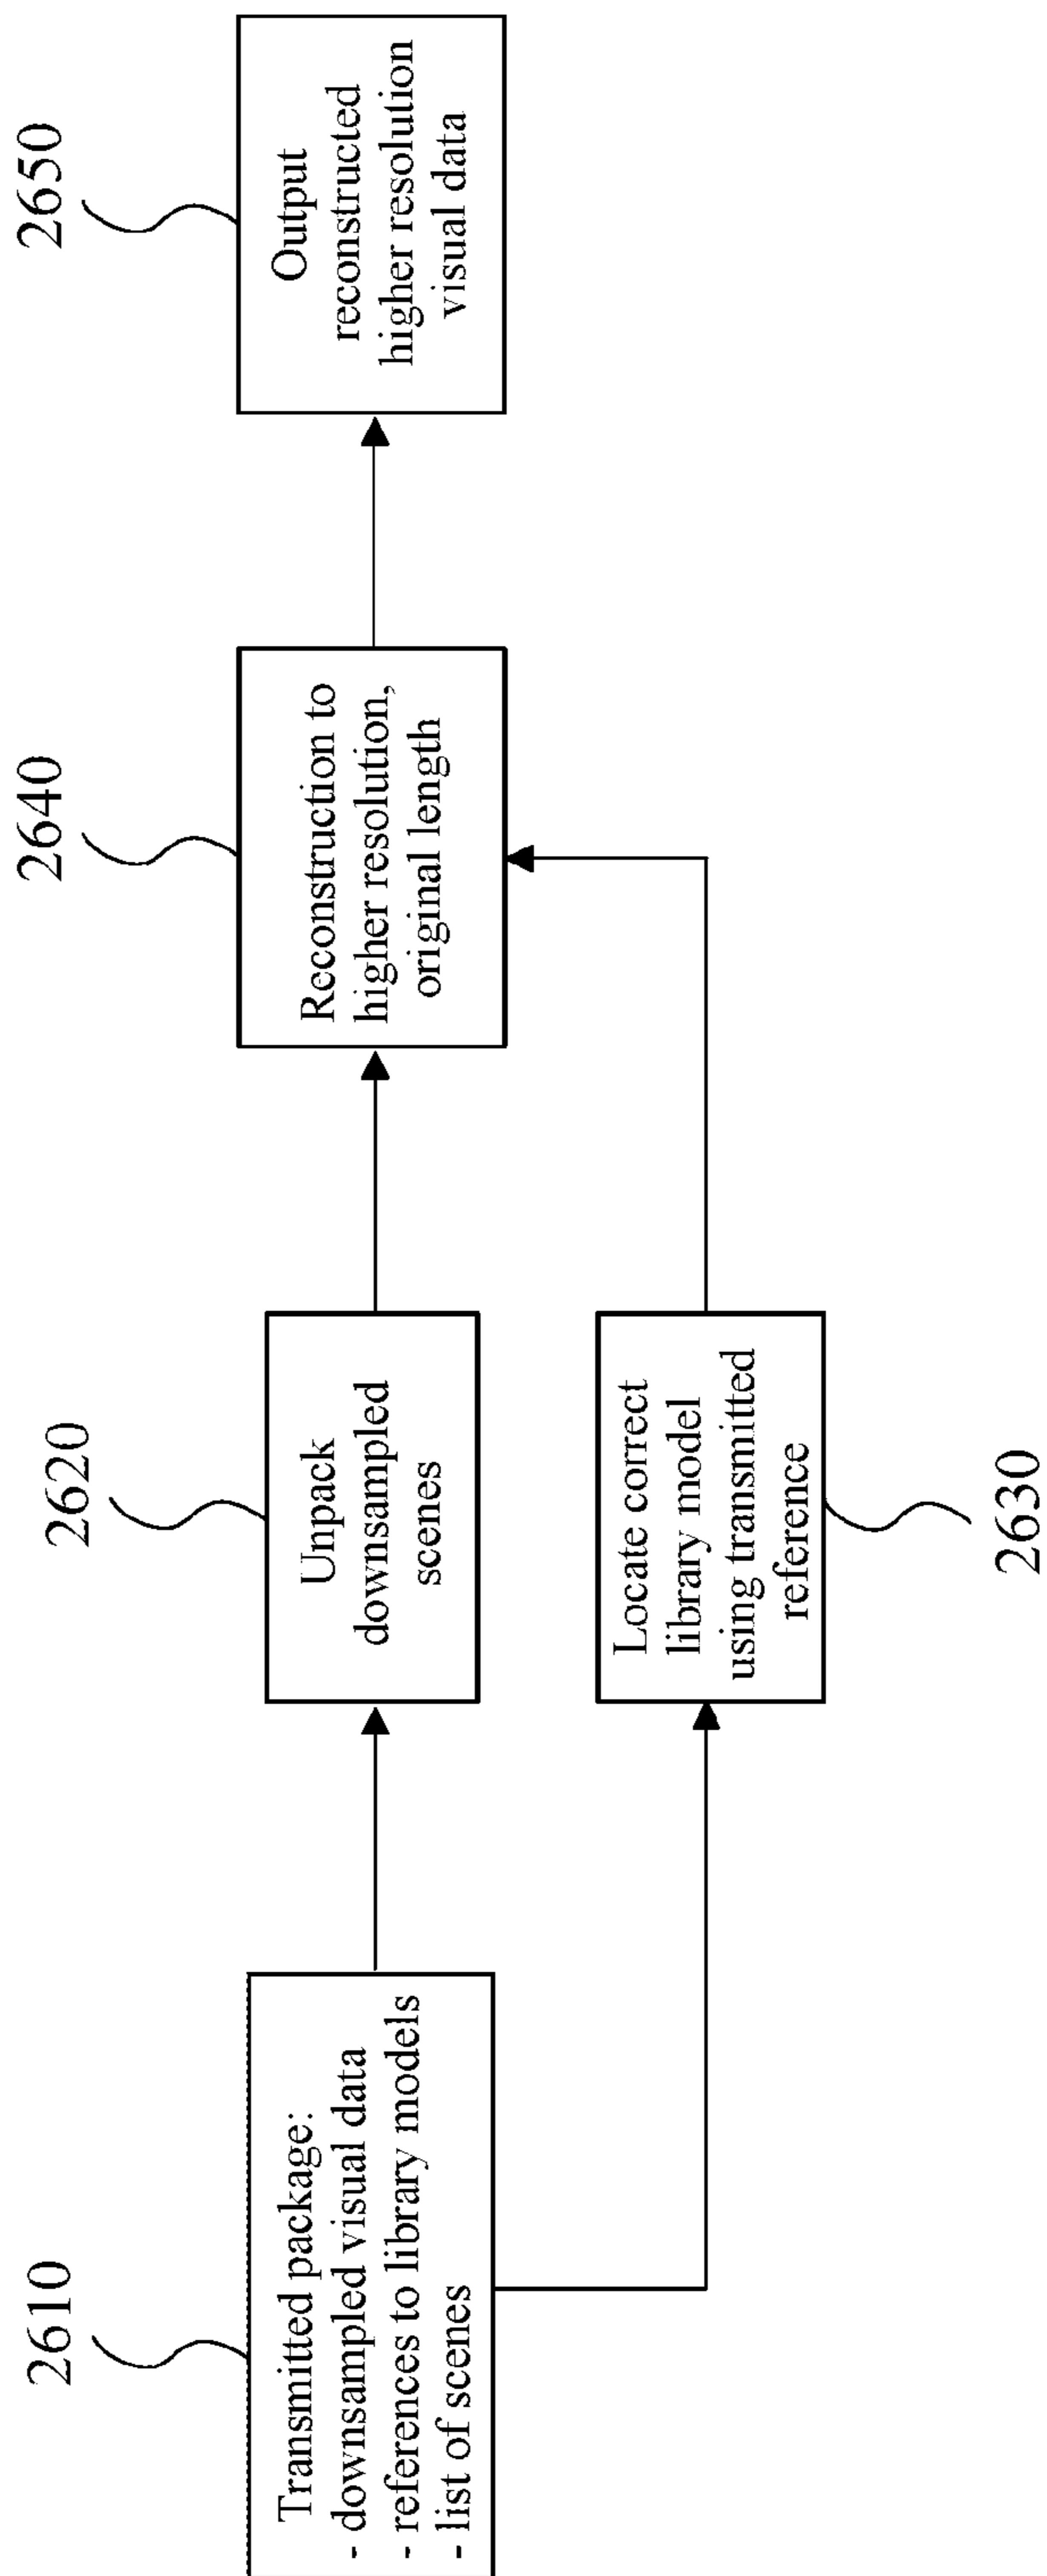



Fig. 27



ACCELERATING MACHINE OPTIMISATION PROCESSES

CROSS REFERENCE TO RELATED APPLICATIONS

This is a continuation of, and claims priority to, International Patent Application No. PCT/GB2016/050427, filed Feb. 19, 2016, which in turn claims priority to the following United Kingdom applications:

GB 1519425.1, filed Nov. 3, 2015;
 GB 1519687.6, filed Nov. 6, 2015;
 GB 1502753.5, filed Feb. 19, 2015;
 GB 1508742.2, filed May 21, 2015;
 GB 1511231.1, filed Jun. 25, 2015;
 GB 1507141.8, filed Apr. 27, 2015;
 GB 1503427.5, filed Feb. 27, 2015; and
 GB 1505544.5, filed Mar. 31, 2015,
 all of which are incorporated herein by reference.

TECHNICAL FIELD

Disclosed implementations relate to methods of developing a learned hierarchical algorithm. Specifically, a hierarchical algorithm is selected based on predetermined metrics and developed before being outputted.

BACKGROUND

Increase in Quality of Video and Display Technology

Developments in display technology have led to significant improvements in the resolution able to be displayed on display hardware, such as on televisions, on computer monitors and using video projectors. For example, television screens that are able to display “High Definition” or “HD” resolution content (typically having a resolution of 1920×1080 pixels) have been broadly adopted by consumers. More recently, television screens able to display Ultra High Definition or “Ultra HD” resolution content (typically having a resolution over 3840×2160 pixels) are starting to become more widespread.

In contrast, HD resolution video content is only now becoming commonplace and most legacy content is only available at either Digital Versatile Disc Video (or “DVD-Video”) resolution (typically having a resolution of 720×586 pixels or 720×480 pixels) or Standard Definition or “SD” resolution (where the video content only has a resolution of 640×480 pixels). Some broadcast channels are limited to SD resolutions. Video-streaming services can be restricted to operating at DVD-Video or SD resolutions, to reduce transmission problems where consumers have limitations on available transmission bandwidth or because of a lack of legacy content at higher resolutions.

As a result, there can be a lack of sufficiently high-resolution video content for display on HD and Ultra HD television screens, for both current video content as well as for legacy video content and video streaming services. Also, over time mobile devices such as mobile ‘phones and tablet computers with increasingly larger and higher-resolution screens are being produced and adopted by users. Further, current video content, being output at HD resolutions, is already at a significantly lower resolution than can be displayed by the latest consumer displays operating at, for example, Ultra HD resolutions. To provide sufficiently immersive virtual reality (or “VR”) experiences, display technology needs to be sufficiently high resolution even for smaller screen sizes.

The user experience of having to display content that has significantly lower resolution than the user’s default screen/display resolution is not optimal.

Growth in Data Transmission and Network Limitations

5 The amount of visual data being communicated over data networks such as the Internet has grown dramatically over time and there is increasing consumer demand for high-resolution, high quality, high fidelity visual data content, such as video streaming including, for example, video at HD and Ultra HD resolution. As a result, there are substantial challenges in meeting this growing consumer demand and high performance video compression is required to enable efficient use of existing network infrastructure and capacity.

10 Video data already makes up a significant fraction of all data traffic communicated over the Internet, and mobile video (i.e. video transmitted to and from mobile devices over wireless data networks such as UTMS/CDMA) is predicted to increase 13-fold between 2014 and 2019, accounting for 72 percent of total mobile data traffic by the end of that forecast period. As a result, there are substantial challenges in meeting this growing consumer demand and more efficient visual data transmission is required to enable efficient use of existing network infrastructure and capacity.

15 To stream video to consumers using available streaming data bandwidth, media content providers can down-sample or transcode the video content for transmission over a network at one or a variety of bitrates so that the resolution of the video can be appropriate for the bitrate available over each connection or to each device and correspondingly the amount of data transferred over the network can be better matched to the available reliable data rates. For example, a significant proportion of current consumer Internet connections are not able to reliably support continuous streaming of video at an Ultra HD resolution, so video needs to be streamed at a lower quality or lower resolution to avoid buffering delays.

20 Further, where a consumer wishes to broadcast or transmit video content, the uplink speeds of consumer Internet connections are typically a fraction of the download speeds and thus only lower quality or lower resolution video can typically be transmitted. In addition, the data transfer speeds of typical consumer wireless networks are another potential bottleneck when streaming video data for video at resolutions higher than HD resolutions or virtual reality data and content to/from contemporary virtual reality devices. A problem with reducing the resolution of a video when transmitting it over a network is that the reduced resolution video may not be at the desired playback resolution, but in some cases there is either not sufficient bandwidth or the bandwidth available is not reliable during peak times for transmission of a video at a high resolution.

25 Alternatively, even without reducing the original video resolution, the original video may have a lower resolution than desired for playback and so may appear at a suboptimal quality when displayed on higher-resolution screens.

Video Compression Techniques

Existing commonly used video compression techniques, such as H.264 and VP8, as well as proposed techniques, such as H.265, HEVC and VP9, all generally use similar approaches and families of compression techniques. These compression techniques make a trade-off between the quality and the bit-rate of video data streams when providing inter-frame and intra-frame compression, but the amount of compression possible is largely dependent on the image resolution of each frame and the complexity of the image sequences.

To illustrate the relationship between bitrate and resolution among other factors, it is possible to use an empirically-derived formula to show how the bitrate of a video encoded with, for example the H.264 compression technique, relates to the resolution of that video:

$$\text{bitrate} \propto Q \times w \times h \times f \times m$$

where Q is the quality constant, w is the width of a video, h is the height of a video, f is the frame-rate of a video and m is the motion rank, where $m \in \{1, \dots, 4\}$ and a higher m is used for fast-changing hard-to-predict content.

The above formula illustrates the direct relationship between the bitrate and the quality constant Q. A typical value, for example, that could be selected for Q would be 0.07 based on published empirical data, but a significant amount of research is directed to optimising a value for Q.

The above formula also illustrates the direct relationship between the bitrate and the complexity of the image sequences, i.e. variable m. The aforementioned existing video codecs focus on spatial and temporal compression techniques. The newer proposed video compression techniques, such as H.265, HEVC and VP9, seek to improve upon the motion prediction and intra-frame compression of previous techniques, i.e. optimising a value for m.

The above formula further illustrates a direct relationship between the bitrate and the resolution of the video, i.e. variables w and h. In order to reduce the resolution of video, several techniques exist to downscale the resolution of video data to reduce the bitrate.

As a result of the disadvantages of current compression approaches, existing network infrastructure and video streaming mechanisms are becoming increasingly inadequate to deliver large volumes of high quality video content to meet ever-growing consumer demands for this type of content. This can be of particular relevance in certain circumstances, for example in relation to live broadcasts, where bandwidth is often limited, and extensive processing and video compression cannot take place at the location of the live broadcast without a significant delay due to inadequate computing resources being available at the location.

Video Upscaling Techniques
To reproduce a video at a higher resolution than that at which it has been transmitted (e.g. by a streaming service or broadcaster) or provided (e.g. on DVD or via a video download provider), various “upsampling” techniques exist to increase the resolution of video data/signals, which enhance image quality when starting from a lower resolution image or video and which produce an image or video of a higher resolution.

Referring to FIG. 14, a conventional upscaling technique 1400 will now be described.

Received video data 1410 is provided into a decoder system and is, for example, a lower-resolution video encoded in a standard video format, such as an SD resolution video. This video format can be a variety of known video codecs, for example such as H.264 or VP8, but can be any video data that the system is able to decode into component frames of video.

The system then separates a first section of the video data 1410 into single frames at step 1420, i.e. into a sequence of images at the full SD resolution of the video data 1410. For some video codecs, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

An upscaling technique 1430 is then used on one or more of the frames or sections of frames, to increase the resolution

of the areas upon which it is used. The higher resolution frames are then optionally processed at step 1440 into a format suitable for output as a video. The video, being composed of higher resolution frames, will be in the form of a higher resolution video 1450 than the original video file.

For example, a basic upscaling technique that makes little attempt to enhance the quality of the video is known as nearest-neighbour interpolation. This technique simply increases the resolution of received video data by representing an original pixel of the transmitted video as multiple pixels or a “block” of pixels. The resulting effect is that the video appears pixelated and in blocks.

Other less basic upscaling techniques use the existing video data to estimate unknown intermediate pixels between known pixel values in order to increase the resolution with a less noticeable loss in quality, these techniques generally known by the term interpolation, these techniques typically factoring into account a weighted average of known pixels in the vicinity of each unknown intermediate pixel or fit to a curve or line to surrounding values and interpolate to the mid-point along the curve or line (e.g. bicubic or bilinear interpolation). Typically, such upscaling techniques determine values for the additional pixels required to create a higher resolution image by averaging neighbouring pixels, which creates a blurring effect or other visual artefacts such as “ringing” artefacts. Most upscaling techniques use interpolation-based techniques to produce higher-resolution versions of received video data. Various methods of interpolation are possible and well documented in the prior art in relation to video or image enhancement.

Various methods of interpolation are possible and well documented in the prior art in relation to video or image enhancement. There are many problems with conventional upscaling techniques. Upscaling techniques that reduce jagged edges tend to introduce more blur to an up-scaled video, for example, while upscaling techniques that reduce “halos” or “ringing” artefacts tend to make an up-scaled video less sharp. Further, conventional upscaling techniques are not content-aware or adaptive. Fundamentally, conventional upscaling techniques are limited by the Nyquist-Shannon sampling theorem.

As a result of the disadvantages of current upscaling techniques, the quality of video data that has been “up-scaled” to a higher resolution than that at which it is stored or transmitted can be inadequate or non-optimal for its intended function.

Super Resolution Techniques for Enhancing Images

Super resolution techniques are techniques that can be described as recovering new high-resolution information that is not explicitly present in low-resolution images.

Super resolution techniques have been developed for many different applications, such as for satellite and for aerial imaging and medical image analysis for example. These applications start with low-resolution images where the higher-resolution image is not available or is possibly unknowable, and by using super resolution techniques it is possible to make substantial enhancements to the resolution of such low-resolution images.

Super resolution techniques allow for the creation of one or more high-resolution images, typically from one or more low-resolution images. Typically, super resolution is applied to a set or series of low-resolution images of the same scene and the technique attempts to reconstruct a higher-resolution image of the same scene from these images.

Super resolution techniques fall predominantly into one of two main fields; optical super resolution techniques and geometrical super resolution techniques. Optical super reso-

lution techniques allow an image to exceed the diffraction limit originally placed on it, while geometrical super resolution techniques increase the resolution from digital imaging sensors. In the field of image resolution enhancement, geometrical super resolution seems to be the predominant technique. Further, super resolution approaches are usually split into learning- or example-based approaches and interpolation-based (multi-frame) approaches. Example based super resolution techniques are generally accepted to be a superior technique to enhance image quality.

One specific super resolution technique is termed multi-exposure image noise reduction. This technique takes the average of many exposures in order to remove unwanted noise from an image and increase the resolution.

Another super resolution technique employed is sub-pixel image localisation, which involves calculating the ‘centre of gravity’ of the distribution of light over several adjacent pixels and correcting blurring accordingly. However, this technique relies on the assumption that all light in the image came from the same source, which is not always a correct assumption.

Machine Learning Techniques

Machine learning is the field of study where a computer or computers learn to perform classes of tasks using the feedback generated from the experience or data gathered that the machine learning process acquires during computer performance of those tasks.

Typically, machine learning can be broadly classed as supervised and unsupervised approaches, although there are particular approaches such as reinforcement learning and semi-supervised learning which have special rules, techniques and/or approaches.

Supervised machine learning is concerned with a computer learning one or more rules or functions to map between example inputs and desired outputs as predetermined by an operator or programmer, usually where a data set containing the inputs is labelled.

Unsupervised learning is concerned with determining a structure for input data, for example when performing pattern recognition, and typically uses unlabelled data sets.

Reinforcement learning is concerned with enabling a computer or computers to interact with a dynamic environment, for example when playing a game or driving a vehicle.

Various hybrids of these categories are possible, such as “semi-supervised” machine learning where a training data set has only been partially labelled.

For unsupervised machine learning, there is a range of possible applications such as, for example, the application of computer vision techniques to image processing or video enhancement. Unsupervised machine learning is typically applied to solve problems where an unknown data structure might be present in the data. As the data is unlabelled, the machine learning process is required to operate to identify implicit relationships between the data for example by deriving a clustering metric based on internally derived information. For example, an unsupervised learning technique can be used to reduce the dimensionality of a data set and attempt to identify and model relationships between clusters in the data set, and can for example generate measures of cluster membership or identify hubs or nodes in or between clusters (for example using a technique referred to as weighted correlation network analysis, which can be applied to high-dimensional data sets, or using k-means clustering to cluster data by a measure of the Euclidean distance between each datum).

Semi-supervised learning is typically applied to solve problems where there is a partially labelled data set, for

example where only a subset of the data is labelled. Semi-supervised machine learning makes use of externally provided labels and objective functions as well as any implicit data relationships.

When initially configuring a machine learning system, particularly when using a supervised machine learning approach, the machine learning algorithm can be provided with some training data or a set of training examples, in which each example is typically a pair of an input signal/vector and a desired output value, label (or classification) or signal. The machine learning algorithm analyses the training data and produces a generalised function that can be used with unseen data sets to produce desired output values or signals for the unseen input vectors/signals. The user needs to decide what type of data is to be used as the training data, and to prepare a representative real-world set of data. The user must however take care to ensure that the training data contains enough information to accurately predict desired output values without providing too many features (which can result in too many dimensions being considered by the machine learning process during training, and could also mean that the machine learning process does not converge to good solutions for all or specific examples). The user must also determine the desired structure of the learned or generalised function, for example whether to use support vector machines or decision trees.

The use of unsupervised or semi-supervised machine learning approaches are sometimes used when labelled data is not readily available, or where the system generates new labelled data from unknown data given some initial seed labels.

Current training approaches for most machine learning algorithms can take significant periods of time, which delays the utility of machine learning approaches and also prevents the use of machine learning techniques in a wider field of potential application.

Machine Learning & Image Super Resolution

To improve the effectiveness of some super resolution techniques, it is possible to incorporate machine learning, otherwise termed a “learned approach”, into the image super resolution techniques described above.

For example, one machine learning approach that can be used for image enhancement, using dictionary representations for images, is a technique generally referred to as dictionary learning. This approach has shown effectiveness in low-level vision tasks like image restoration.

When using dictionary learning, the representation of a signal is given as a linear combination of functions drawn from a collection of atoms referred to as a dictionary. For example, a given signal y can be represented as:

$$y = \alpha_1 x_1 + \alpha_2 x_2 + \dots + \alpha_n x_n$$

where x_1, \dots, x_n are the atoms of a dictionary of size n and $\alpha_1, \dots, \alpha_n$ are coefficients such that $\|\alpha\|_0 < \lambda$, where λ is the sparsity constraint, for example where $\lambda=3$ no more than three coefficients can be non-zero. The atoms have the same dimensionality as the signal y so, while it is possible to have an atom x_i that is identical to y , a dictionary of simple atoms can usually be used to reconstruct a wide range of different signals.

In theory, at least k orthogonal atoms are required to fully reconstruct signals in k -dimensional space. In practice, however, improved results are achieved through using an over-complete dictionary where there are $n > k$ atoms and these atoms do not have to be orthogonal to one another.

A complete dictionary means that the number of dictionary atoms is the same as the dimensionality of the image

patches and that the dictionary atoms are linearly independent (i.e. all orthogonal to each other and can represent the entire, or complete, dimensional space), so where 16×16 atoms represent 16×16 image patches, the dictionary is complete if it has $16 \times 16 = 256$ atoms. If more atoms than this are present in the dictionary, then the dictionary becomes over-complete.

An example of an over-complete dictionary is shown in FIG. 1, where a 16×16 pixel patch is represented by a linear combination of 16×16 dictionary atoms **5** that is drawn from the collection of atoms that is the dictionary **1**. It is noted that the atoms are not selected locally within the dictionary, but instead are chosen as the linear combination that best approximates the signal patch for a maximum number of atoms allowed and irrespective of their location within the dictionary. Without a constraint that the atoms must be orthogonal to one another, larger dictionaries than the signal space that the dictionary is intended to represent are created.

Over-complete dictionaries are used because they provide better reconstructions, but at the cost of needing to store and transmit all of the new dictionaries and representations created during the dictionary learning process. In comparison with a predetermined library of representations, a significantly increased amount of data is created as a result of dictionary learning because it generates a data set significantly larger than the basis set in a predetermined library of representations and the atoms are not all orthogonal to one another.

In dictionary learning, where sufficient representations are not available in an existing library of representations (or there is no library available), machine learning techniques are employed to tailor dictionary atoms such that they can adapt to the image features and obtain more accurate representations. Each new representation is then transferred along with the video data to enable the representation to be used when recreating the video for viewing.

The transform domain can be a dictionary of image atoms, which can be learnt through a training process known as dictionary learning that tries to discover the correspondence between low-resolution and high-resolution sections of images (or “patches”). Dictionary learning uses a set of linear representations to represent an image and, where an over-complete dictionary is used, a plurality of linear representations can be used to represent each image patch to increase the accuracy of the representation.

When using dictionary learning based super resolution techniques, there is a need for two dictionaries: one for the low-resolution image and a separate dictionary for the high-resolution image. To combine super resolution techniques with dictionary learning, reconstruction models are created to enhance the image based on mapping the coefficients of the low-resolution dictionary to coefficients in the high-resolution dictionary. Various papers describe this, including “On Single Image Scale-Up Using Sparse-Representations” by R. Zeyde et al and published in 2010, “Image super-resolution via sparse representation” by J. Yang and published in 2010, and “Coupled Dictionary Training for Image Super-Resolution” by J. Yang et al and published in 2012, which are incorporated by reference.

A disadvantage of using dictionary learning based super resolution techniques on low-resolution images to attempt to recreate the high-resolution image is the need for two dictionaries, one for the low-resolution image and a separate dictionary for the high-resolution image. It is possible to have a single combined dictionary, but in essence there is

always in practice an explicit modelling for each resolution to enable representations to be matched between the two resolutions of image.

A further disadvantage of using dictionary learning, however, especially when used with an over-complete dictionary, is the amount of data that needs to be transferred along with the low-resolution image in order to recreate a high-resolution image from the low-resolution image.

Another disadvantage of dictionary learning approaches is that these tend to use a local patch averaging approach in the final step of reconstruction of a higher-resolution image from a lower-resolution image, which can result in unintentional smoothing in the reconstructed image.

Another further disadvantage of dictionary learning approaches is that it is very slow and can have high memory requirements, depending on the size of the dictionary.

Artefact Removal in Visual Data

Visual data artefacts and/or noise can often be introduced into visual data during processing, particularly during processing to compress visual data or during transmission of the visual data across a network. Such introduced artefacts can include blurring, pixelation, blocking, ringing, aliasing, missing data, and other marks, blemishes, defects, and abnormalities in the visual data. These artefacts in visual data can degrade the user experience when viewing the visual data. Furthermore, these artefacts in visual data can also reduce the effectiveness of visual data processing techniques, such as image super resolution, as well as other visual tasks such as image classification and segmentation, that use processed images as an input.

Lossy compression, in which a visual data is encoded using inexact approximations, is a particularly common source of artefacts. Lossy compression is often required to reduce the size of a digital image or video in order to transmit it across a network without using an excessive amount of bandwidth. High visual data compression ratios can be achieved using lossy compression, but at the cost of a reduction in quality of the original visual data and the introduction of artefacts.

The transmission of visual data across a network can itself introduce artefacts to the visual data through transmission errors between nodes in the network.

Current methods of removing artefacts (herein referred to as visual data fidelity correction) generally only correct a specific types of artefact. Examples of such techniques include deblocking oriented methods, such as Pointwise Shape-Adaptive Discrete Cosine Transform (hereafter referred to as Pointwise SA-DCT), which deal with blocking artefacts. Such techniques do not perform well, and can also introduce further artefacts as a side effect, such as the over smoothing of visual data.

SUMMARY

Aspects and/or embodiments are set out in the appended claims. These and other aspects and embodiments are also described herein.

Certain aspects and/or embodiments seek to provide techniques for generating hierarchical algorithms that can be used, when converting original high-quality visual data into lower-quality visual data, to allow recreation of higher-quality visual data from the lower-quality visual data without significant loss in quality between the original high-quality visual data and the higher quality visual data.

Further aspects and/or embodiments seek to provide techniques for reconstruction and/or enhancement from lower-quality visual data to higher-quality visual data.

Other aspects and/or embodiments seek to provide techniques for machine learning.

General Training Method

According to an aspect, there is provided a method for developing an enhancement model for low-quality visual data, the method comprising the steps of: receiving one or more sections of higher-quality visual data; training a hierarchical algorithm, wherein the hierarchical algorithm is operable to increase the quality of one or more sections of lower-quality visual data to substantially reproduce the one or more sections of higher-quality visual data; and outputting the hierarchical algorithm.

Training hierarchical algorithms can allow enhancement and/or reconstruction models to be developed for enhancing visual data in at least one embodiment. In some embodiments visual data can be image data and/or video data. Furthermore, in some embodiments enhancement models can be developed to increase the accuracy of higher-quality visual data that can be reproduced from lower-quality visual data in comparison to original higher-quality visual data. In at least one embodiment, knowledge of the original visual data can allow the hierarchical algorithm to be trained (and/or developed) based on knowledge of both the original visual data and the low-quality visual data in order to train a hierarchical algorithm to substantially reproduce the original visual data from the low-quality visual data.

Optionally, the hierarchical algorithm is developed from a known initialisation.

In some embodiments, the hierarchical algorithm can be developed from a known initialisation, for example of a hierarchical function or basis. In some of these embodiments, the hierarchical function or basis can be for example haar wavelets or one or more pre-trained hierarchical algorithms or sets of hierarchical algorithms. In at least one embodiment, providing a known initialisation allows the training of hierarchical algorithms to be accelerated, and the known initialisation can be closer to the best solution especially when compared to starting from a random initialisation. In some embodiments, a trained hierarchical algorithm can be developed for input visual data, wherein the trained hierarchical algorithm is developed for that input data based on the selected most similar pre-trained algorithm. In at least one embodiment, the selection of the one or more most similar pre-trained algorithm(s) can be made based on one or more metrics associated with the pre-trained models when compared and/or applied to the input data. In some embodiments, metrics can be any predetermined measure of similarity or difference. In some embodiments, the most similar pre-trained algorithm can be used as a starting point for developing a trained or tailored algorithm for the input data as a tailored algorithm does not have to undergo as extensive development as needed when developing an algorithm from first principles.

Optionally, the method is performed at a first network node within a network. Furthermore, optionally the hierarchical algorithm may be transmitted to a second network node in the network.

By lowering the quality of visual data (for example by lowering the resolution of video data) in some embodiments, less data can be sent across a network from a first node to a second node in order for the second node to display the visual data from the first node. In some embodiments, the lower quality visual data together with a model to be used for reconstruction can allow for less data to be transmitted than if the original higher-quality version of the same visual data is transmitted between nodes.

Optionally, the one or more sections of lower-quality visual data are generated from the one or more sections of higher-quality visual data. Furthermore, optionally the one or more sections of lower-quality visual data may be generated from the high-quality visual data using a process comprising down-sampling.

By lowering the quality of a section of visual data in some embodiments, less data can be sent in order to transmit the visual data over a network. Further, in some embodiments, sending both the lower quality version together with a model to be used for reconstruction can still result in less data being transmitted than if the original higher-quality version of the same section of visual data is transmitted alone.

Optionally, the one or more sections of lower-quality visual data are generated from the one or more sections of higher-quality visual data using a process comprising compression and/or quantisation.

Lossy compression, a reduction in frame rate, a reduction in pixel data precision (e.g. from 32-bit to 16-bit) and quantisation of visual data are methods for producing lower-quality visual data from higher-quality visual data and can be used in some embodiments to generate lower-quality visual data from higher-quality visual data.

Optionally, the one or more sections of higher-quality visual data and/or lower quality visual data comprise any of: a single frame, a sequence of frames and a region within a frame or sequence of frames. Optionally, the one or more sections of higher-quality visual data and/or the one or more sections of lower-quality visual data may comprise an image, a sequence of images, a section of video or an output of a video game.

Depending on the visual data being processed in an embodiment, models can be generated for sections of visual data comprising a single frame, a sequence of frames or a region within a frame or sequence of frames. Each of these options can be used in some or all embodiments in order to provide a method of enhancing or reconstructing visual data to produce higher quality visual data.

General Enhancement/Reconstruction

According to another aspect, there is provided a method for enhancing lower-quality visual data using hierarchical algorithms, the method comprising the steps of: receiving one or more sections of lower-quality visual data; applying a hierarchical algorithm to the one or more sections of lower-quality visual data to enhance the one or more sections of lower-quality visual data to one or more sections of higher-quality visual data, wherein the hierarchical algorithm was developed using a learned approach; and outputting the one or more sections of higher-quality visual data.

In some embodiments, a section of visual data that has been transmitted across a network can be enhanced using hierarchical algorithms. By applying a hierarchical algorithm in some or all of these embodiments, a higher-quality version of the visual data can be output for input lower-quality visual data. Therefore, in some embodiments, only lower-quality visual data need to be transmitted over the network. In other embodiments, lower-quality visual data can be transmitted over the network along with one or more hierarchical algorithms that can be used to enhance lower-quality visual data.

Optionally, the hierarchical algorithm is selected from a library of learned hierarchical algorithms.

In some embodiments, a stored library of learned hierarchical algorithms allows selection of a hierarchical algorithm for comparison without having to develop them or obtain them from an external source. In some embodiments, the comparison can be between a plurality of algorithms in

the library. Use of such a library, in at least some embodiments, may result in the faster selection of a suitable hierarchical algorithm for enhancing the visual data or, in some embodiments, the most suitable hierarchical algorithm in a library (for example, by basing a measure of suitability on a predetermined metric).

Optionally, the selection of the hierarchical algorithm from the library of learned hierarchical algorithms is determined by metric data associated with the lower-quality visual data.

In some embodiments, it is assumed that the closer the features of the hierarchical algorithm are to those of metric data associated with the lower-quality visual data, the more accurate a reconstruction can be created using that particular hierarchical algorithm. Therefore, by using associated metric data in some embodiments in this way, an appropriate model may be more accurately chosen from the plurality of hierarchical algorithms available.

Optionally, the steps of receiving the one or more sections of lower-quality visual data and applying the hierarchical algorithm to the one or more sections of lower quality visual data occur substantially simultaneously.

In some embodiments, by receiving the one or more sections of lower-quality visual data and substantially simultaneously applying the hierarchical algorithm to one or more sections of lower-quality visual data, the time taken to enhance the visual data can be reduced. This is especially beneficial for live broadcasting embodiments, where it can be advantageous for the time taken for visual data processing before transmission to be minimised.

Optionally, the one or more sections of lower-quality visual data are generated from one or more sections of original higher-quality visual data.

In some embodiments, by generating lower-quality visual data from original higher-quality visual data and transmitting the lower-quality visual data instead of the original higher-quality visual data, less data needs to be processed or sent across a network.

Offline Training

According to another aspect, there is provided a method for developing an enhancement model for visual data, the method comprising the steps of: receiving one or more sections of higher-quality visual data and one or more sections of lower-quality visual data corresponding to the one or more sections of higher-quality data; developing a hierarchical algorithm, wherein the hierarchical algorithm is operable to substantially reproduce the one or more sections of higher-quality visual data from the corresponding one or more sections of lower-quality visual data; and storing the developed hierarchical algorithm in a library of developed hierarchical algorithms.

In some embodiments, by developing and storing hierarchical algorithms trained to enhance visual data in a library of enhancement algorithms, trained hierarchical algorithms can be selected from the library for repeated use to reduce the need for individual training from scratch every time enhancement of visual data is required. In some of these embodiments, this approach can save computational expense and speed up the transmission of visual data across a network, for example where it is possible to choose a suitable model for visual enhancement from the library rather than train a new one. It also allows for hierarchical algorithms to be trained in advance of the transmission of visual data, for at least some embodiments, removing time constraints that can be present when training hierarchical algorithms substantially immediately before visual data is transmitted.

Optionally, the developed hierarchical algorithm is stored associated with metric data relating to a content type of the higher-quality visual data and/or lower-quality visual data from which the hierarchical algorithm was developed.

In some embodiments, by storing the trained model in the library with metric data relating to the associated visual data, optionally where the metric data can be used to determine whether the trained model would be suitable for use with visual data having similar metric data, such metric data can be used to compare to metric data associated with further visual data to select a suitable hierarchical algorithm to enhance that further visual data. This means that, in at least some of these embodiments, a specific hierarchical algorithm need not be trained for every set of visual data; existing trained hierarchical algorithms can instead be used to enhance similar sets of visual data to that on which they were trained.

Optionally, the hierarchical algorithm is developed from a known hierarchical algorithm.

In some embodiments, developing a new hierarchical algorithm from a known hierarchical algorithm that was trained on similar visual data to the visual data on which the new algorithm is to be trained can reduce the time and/or computational effort required to train the new hierarchical algorithm.

Optionally, the known hierarchical algorithm is stored in the library of developed hierarchical algorithms.

For some embodiments, having one or more known algorithms stored in a library of previously developed algorithms can allow for these algorithms to be accessed quickly and easily, increasing the efficiency of the training process.

Optionally, the method is performed at a network node within a network.

By performing the method at a network node in some embodiments, the resulting trained hierarchical algorithms can be transmitted across the network to further network nodes at which they are required. It also allows, in at least some embodiments, for the library of trained models to be stored remotely from the location at which the training takes place.

Optionally, the method additionally comprises a step of encoding the one or more sections of higher quality visual data and/or the one or more sections of lower-quality visual data.

Optionally, the step of encoding the one or more sections of higher quality visual data and/or the one or more sections of lower-quality visual data occurs after the step of developing the hierarchical algorithm.

Visual data is often encoded prior to transmission across a network, or storage in a memory, in embodiments. By encoding the visual data after the hierarchical algorithm has been trained in at least some embodiments, the visual data is then ready for transmission or storage. The trained hierarchical algorithm will then be operable to enhance the quality of decoded visual data in these embodiments.

Optionally, the method additionally comprises a step of refining the hierarchical algorithm after the step of encoding the one or more sections of higher quality visual data and/or the one or more sections of lower-quality visual data.

By refining the hierarchical algorithm after the encoding of the visual data in some embodiments, the hierarchical algorithm can be trained to enhance the quality of the encoded visual data without the visual data having to be decoded first. In some of these embodiments, this can decrease the time required to decode and enhance the lower-quality visual data to substantially reproduce the higher-quality visual data.

According to a further aspect, there is provided a method for reducing the amount of data to be transferred when communicating visual data over a network from a first node to a second node, the method at the first node comprising the steps of: reducing the quality of one or more sections of higher-quality visual data to one or more sections of lower-quality visual data; selecting at least one algorithm operable to increase the quality of at least one section of lower-quality visual data using the higher-quality visual data to optimise the selection of the algorithm, wherein the algorithm corresponds to at least the one section of lower quality visual data; transmitting the one or more sections of lower-quality visual data to the second node; and transmitting to the second node one or more references corresponding to the one or more selected algorithms that correspond to the one or more sections of lower-quality visual data transmitted to the second node; wherein the second node is able to identify the selected algorithm using the transmitted one or more references and substantially reproduce the higher-quality visual data from the transmitted lower-quality visual data using the identified selected algorithms that corresponds to each section of lower-quality visual data.

According to another aspect, there is provided a method for increasing the quality of a section of visual data communicated over a network from a first node to a second node, the method at the second node comprising the steps of: receiving a lower-quality visual data via a network; receiving a corresponding reference to an algorithm operable to increase the quality of the lower-quality visual data, the algorithm selected with the knowledge of a higher-quality visual data from which the lower-quality visual data was generated; and using the algorithm to increase the quality of the lower-quality visual data to substantially recreate the higher-quality visual data.

According to a further aspect, there is provided a system for reducing the amount of data transferred when communicating visual data over a network, the system comprising two or more nodes wherein a first node is configured to: reduce the quality of one or more sections of higher-quality visual data to one or more sections of lower-quality visual data; select at least one algorithm operable to increase the quality of at least one section of lower-quality visual data using the higher-quality visual data to optimise the selection of the algorithm, wherein the algorithm corresponds to at least the one section of lower quality visual data; transmit the one or more sections of lower-quality visual data to the second node; and transmit to the second node one or more references to the one or more algorithms that correspond to the one or more sections of lower-quality visual data transmitted to the second node; wherein the second node is configured to: receive a lower-quality visual data via a network; receive a corresponding reference to the algorithm operable to increase the quality of the lower-quality visual data, the algorithm selected with the knowledge of a higher-quality visual data from which the lower-quality visual data was generated; and use the algorithm to increase the quality of the lower-quality visual data to substantially recreate the higher-quality visual data.

By transmitting a lower-quality version of a visual data in some embodiments, such as a section of low-quality visual data or series of visual data sections, together with a library reference to an algorithm (i.e. any or all of an algorithm, reconstruction algorithm, model, reconstruction model, parameters or reconstruction parameters) to aid reconstruction of higher quality visual data, such as a high-resolution video frame or series of frames, in at least some embodiments less data can be transferred over a network to enable

high-quality visual data to be viewed compared to transmitting the high quality visual data alone.

Optionally, the steps of transmitting the one or more sections of lower-quality visual data to the second node and transmitting to the second node one or more references corresponding to the one or more selected algorithms that correspond to the one or more sections of lower-quality visual data transmitted to the second node occur together, or substantially simultaneously.

By transmitting both visual data and one or more references to algorithms in a library of algorithms, a reduced amount of data can be transmitted as only one or more references to algorithms are transmitted instead of the algorithms themselves.

Optionally, the algorithm is a hierarchical algorithm.

In some embodiments, the algorithms used are hierarchical algorithms. It should be noted that algorithms could also be referred to as models, representations, parameters or functions. In some of these embodiments, hierarchical algorithms can enable substantially accurate reconstruction of visual data, e.g. produce a higher quality high-resolution video from the low-resolution video that is transmitted, for example where quality can be measured by a low error rate in comparison to the original high-resolution video.

Optionally, the algorithm is a non-linear algorithm.

In some embodiments, the use of non-linear algorithms can be more flexible and expressive than dictionary learning based approaches, and use fewer coefficients for the reconstruction of higher-quality visual data. In some of these embodiments, this can allow the reconstruction of the sections of higher-quality to be substantially accurate.

Optionally, the algorithm is selected from a library of algorithms stored at any of: the first node; the second node; a centralised database in the network; or a distributed database in the network.

In some embodiments, a library of algorithms can allow for the selection of an substantially optimal, if not the most optimal, algorithm available in the library to reconstruct the lower-quality visual data into higher quality visual data. In some of these embodiments, the use of a library can also allow the selected algorithm to be referred to by a reference identifier. In certain embodiments, libraries can be provided at both nodes, and/or in centralised or distributed databases, and optionally can use common or synchronised reference identifiers for the same algorithms.

Optionally, the received reference corresponds to an algorithm stored in a library at the second node.

By providing common or synchronised libraries of algorithms at both the first and the second nodes in at least some embodiments, and by transmitting a reference or reference identifier when transmitting the corresponding lower-quality visual data to allow selection of matching algorithms from both libraries using the reference identifier to identify the selected algorithm, only the reference identifier and the lower-quality visual data needs to be transmitted between the nodes thus data transmission is reduced as the algorithm itself doesn't need to be transmitted.

Optionally, if the second node cannot identify the selected algorithm, the second node sends a request to any of: the first node; a centralised database; or a distributed database for transmission of the selected algorithm to the second node.

In some embodiments, by configuring the second node to be able to request models from a node (for example a first node or alternatively another node or a node to which multiple libraries are synchronised, depending on embodiment) in situations where the libraries at the second node and other node are not synchronised, the higher-quality visual

data can still be reconstructed even if the transmitted reference does not correspond to an algorithm stored at the second node. This can prevent errors in the reconstruction process in some embodiments.

Optionally, the algorithm comprises one or more convolutional neural networks.

In some embodiments, convolutional neural networks can achieve substantially optimal reconstruction of higher-quality visual data from lower-quality visual data (when the reconstructed higher-quality visual data is compared to the original high-quality visual data) while being relatively small in data size, for example when compared to over-complete dictionaries.

Optionally, the high-quality visual data is divided into smaller sections based on similarities between a plurality of sections of the high-quality visual data.

In some embodiments, by dividing the visual data into smaller sections, where the sections can be sequences of frames or portions of one or more frames, and where the division can be based on a particular metric for similarity, more efficient models can be selected. For example, in some embodiments multiple sections can be grouped, all of which comprise part of a landscape shot, and one model can be used to reconstruct the scene, i.e. sequence of frames, as opposed to a using a different model for every separate frame in the scene. In some embodiments, if the next scene in the visual data is very different (for example a scene with significant movement after a scene showing a still landscape), then the scene can be detected as being very different and a new model can be selected accordingly for the scene. In some embodiments, specific models can be selected for each scene or section, allowing at least some optimisation or adapting of the reconstruction model(s) in comparison to the use of a generic model for the whole of the visual data.

Optionally the one or more sections comprises any of: a single frame, a sequence of frames and a region within a frame or sequence of frames.

Depending on the visual data being processed, in some embodiments models can be selected for sections of visual data comprising a single frame, a sequence of frames or a region within a frame or sequence of frames. This selection can be necessary in some of these embodiments in order to provide the most efficient method of reconstructing a higher quality version of a part of the original visual data.

Optionally, the visual data is converted into a sequence of sections, optionally before the quality is lowered.

Separating the visual data into a sections, for example into a series of frames or images in some embodiments, allows for the individual sections to be down-sampled, thus reducing the size of visual data, and thereby can allow for lower quality sections to be transmitted as re-encoded visual data in the original code (or optionally re-encoded using a more optimal codec and/or at a lower quality in some embodiments).

Optionally, down-sampling is used to reduce the quality of one or more sections of higher-quality visual data to one or more sections of lower-quality visual data.

By lowering the quality of the section of visual data in some embodiments, less data needs to be sent across a network from a first node to a second node in order for the second node to receive the file. In these embodiments the lower quality version together with a reference to the model to be used for reconstruction can still allow for less data to be transmitted than if the original higher-quality version of the same section of visual data is transmitted.

Optionally, the steps of transmitting the one or more sections of lower-quality visual data to the second node and

transmitting to the second node the references to the one or more algorithms that correspond to the one or more sections of lower-quality visual data transmitted to the second node occur substantially simultaneously.

By transmitting the lower quality section of visual data and the model reference simultaneously in some embodiments, the reconstruction process can be initiated substantially instantaneously on arrival of the data, as opposed to there being a delay between the arrival of the section of visual data and the model reference or vice versa. In some embodiments, both the model and the section of visual data are necessary for a higher-quality version of the section of visual data to be reconstructed such that the reconstructed visual data is of substantially the same quality as the original higher-quality visual data.

Online Training

According to an aspect, there is provided a method for enhancing visual data when communicating visual data over a network from a first node to a second node, the method at the first node comprising the steps of: reducing the quality of one or more sections of higher-quality visual data to one or more sections of lower-quality visual data; developing at least one hierarchical algorithm operable to increase the quality of the one or more sections of lower quality visual data using the one or more sections of higher-quality visual data to enhance the developed at least one hierarchical algorithm, wherein the developed at least one hierarchical algorithm corresponds to the one or more sections of lower quality visual data; transmitting the one or more sections of lower-quality visual data to the second node; and communicating to the second node at least one of the developed at least one hierarchical algorithms that corresponds to the one or more sections of lower-quality visual data transmitted to the second node; wherein the second node is able to substantially reproduce the one or more sections of higher-quality visual data from the transmitted one or more sections of lower-quality visual data using the developed at least one hierarchical algorithm that corresponds to the one or more sections of lower-quality visual data.

According to a further aspect, there is provided a method for enhancing visual data when communicating visual data over a network from a first node to a second node, the method at the second node comprising the steps of: receiving one or more sections of lower-quality visual data via a network; receiving a corresponding at least one hierarchical algorithm operable to increase the quality of the one or more sections of lower-quality visual data, the at least one hierarchical algorithm developed with the knowledge of one or more sections of higher-quality visual data corresponding to the one or more sections of lower-quality visual data; and using the at least one hierarchical algorithm to increase the quality of the one or more sections of lower-quality visual data to substantially recreate the one or more sections of higher-quality visual data.

According to another aspect, there is provided a system for reducing the amount of data transferred when communicating visual data over a network, the system comprising two or more nodes wherein a first node is configured to: reduce the quality of one or more sections of higher-quality visual data to one or more sections of lower-quality visual data; develop at least one hierarchical algorithm operable to increase the quality of the one or more sections of lower quality visual data using the one or more sections of higher-quality visual data to enhance the developed at least one hierarchical algorithm, wherein the developed at least one hierarchical algorithm corresponds to the one or more sections of lower quality visual data; transmitting the one or

more sections of lower-quality visual data to the second node; and communicate to the second node at least one of the developed at least one hierarchical algorithms that corresponds to the one or more sections of lower-quality visual data transmitted to the second node; wherein the second node is configured to: receive one or more sections of lower-quality visual data via a network; receive a corresponding at least one hierarchical algorithm operable to increase the quality of the one or more sections of lower-quality visual data, the at least one hierarchical algorithm developed with the knowledge of one or more sections of higher-quality visual data corresponding to the one or more sections of lower-quality visual data; and use the at least one hierarchical algorithm to increase the quality of the one or more sections of lower-quality visual data to substantially recreate the one or more sections of higher-quality visual data.

In some embodiments, by developing a hierarchical algorithm for enhancing the quality of a specific set of lower-quality visual data, a more accurate enhancement of that visual data can be achieved at a receiving node in a network. In some of these embodiments, having knowledge of the higher-quality visual data to which the lower-quality visual data corresponds can lead to the training and/or development of more accurate hierarchical algorithms. In further of these embodiments, having knowledge of the compression settings of the lower-quality visual data can also help in training and/or development of more accurate hierarchical algorithms and, in such embodiments, training hierarchical algorithms on visual data having the same compression settings can provide substantially accurate hierarchical algorithms for those compression settings. Furthermore, in some embodiments, communicating the lower-quality algorithm across the network along with the associated lower-quality visual data can reduce the amount of data required to be transmitted across the network when compared with the transmission of the higher-quality data alone. For some embodiments, this is especially true in the case where the higher-quality visual data is at a significantly higher quality than the lower-quality visual data.

Optionally, the developed at least one hierarchical algorithm is selected from a plurality of hierarchical algorithms developed in parallel at the first node.

Optionally, the at least one hierarchical algorithm is developed in parallel to the encoding of the lower-quality visual data.

In some embodiments, developing multiple hierarchical algorithms in parallel while encoding the visual data can speed up the process of preparing the visual data and hierarchical algorithm for communication across a network. In some of these embodiments, by developing multiple hierarchical algorithms in parallel, a greater number of possible algorithm structures can be explored and the most suitable chosen.

Optionally, the steps of transmitting the one or more sections of lower-quality visual data to the second node and communicating to the second node the developed at least one hierarchical algorithm that corresponds to the one or more sections of lower-quality visual data transmitted to the second node occur substantially simultaneously.

By transmitting the visual data and communicating the associated hierarchical algorithm substantially simultaneously in some embodiments, the reconstruction process at the second node can begin as soon as the hierarchical algorithm is received. This can reduce the delay, in some of these embodiments, between receiving lower-quality visual data and outputting higher-quality visual data.

Optionally, the communicating to the second node of the at least one of the developed at least one hierarchical comprises transmitting the developed at least one hierarchical algorithm to the second node.

In some embodiments, transmitting the whole of the developed hierarchical algorithm to the second node ensures that the developed hierarchical algorithm is available for use at the second node.

According to a further aspect, there is provided a method for reducing the amount of data to be transferred when communicating visual data over a network from a first node to a second node, the method at the first node comprising the steps of: reducing the quality of a section of higher-quality visual data to a section of lower-quality visual data; analysing a first example based model for the use of reconstruction of the lower-quality section of visual data; analysing a further one or more example based models for the use of reconstruction of the lower-quality section of visual data; selecting one or more example based models from a plurality of example based models to use based on a specified metric; transmitting the section of lower-quality visual data to the second node; and transmitting to the second node the one or more selected example based models that correspond to the lower-quality section of visual data transmitted to the second node; wherein the second node is able to substantially reproduce the higher-quality section of visual data from the transmitted lower-quality section of visual data using the one or more example based models that corresponds to the lower-quality section of visual data.

According to another aspect, there is provided a method for increasing the quality of a lower-quality section of visual data communicated over a network from a first node to a second node, the method at the second node comprising the steps of: receiving a lower-quality section of visual data via a network; receiving a corresponding example based model operable to increase the quality of the lower-quality section of visual data; and using the example based model to increase the quality of the lower-quality section of visual data to substantially recreate the higher-quality section of visual data.

According to a further aspect, there is provided a system for reducing the amount of data to be transferred when communicating visual data over a network, the system comprising two or more nodes wherein a first node is configured to: reduce the quality of a higher-quality section of visual data to a lower-quality section of visual data; analyse a first example based model for the use of reconstruction of the lower-quality section of visual data; analyse one or more example based models for the use of reconstruction of the lower-quality section of visual data; select one or more example based models from a plurality of example based models to use based on a specified metric; transmit the section of lower-quality visual data to the second node; and transmit to the second node the one or more example based models that correspond to the section of lower-quality visual data transmitted to the second node; wherein the second node is configured to: receive a lower-quality section of visual data via network; receive one or more corresponding example based models operable to increase the quality of the lower-quality section of visual data; and use the one or more example based models to increase the quality of the lower-quality section of visual data to substantially recreate the higher-quality section of visual data.

According to an aspect, there is provided a method for increasing the quality of a lower-quality section of visual data, the method at the first node comprising the steps of:

analysing a first example based model for the use of increasing the quality of the lower-quality section of visual data; analysing a further one or more example based models for the use of increasing the quality of the lower-quality section of visual data; and selecting one or more example based models from a plurality of example based models to use based on a specified metric; wherein the quality of the lower-quality section of visual data is able to be substantially increased using the one or more example based models that corresponds to the lower-quality visual data.

In some embodiments, by transmitting a section of lower-quality visual data over a network together with an example based model to aid reconstruction of high-quality visual data, less data can be transferred over the network to enable a higher-quality visual data to be displayed when compared to transmitting higher-quality visual data alone. In some of these embodiments, by selecting one or more example based models from a plurality of example based models, higher-quality visual data is operable to be reconstructed from the lower-quality visual data and the one or more example based models together. As any example based models used in such embodiments are already in existence and just need to be selected and not created, the delay between the recording of the higher-quality visual data and emission of the reconstructed higher-quality visual data can be minimised. Some of these embodiments can therefore be used for the live broadcast of visual data pertaining to an event. As there exists a plurality of example based models, in some embodiments optional analysis of a plurality of models allows for the selection of an example based model which, when used in the reconstruction of a lower-quality section of visual data, will generate the highest quality output data on reconstruction compared to other selectable models. The visual data to be transmitted can be recorded in the form of higher-quality visual data in some embodiments, the quality of which is then reduced before transmission as outlined above. The visual data can, in alternative embodiments, be recorded in the form of lower-quality visual data, which requires no further decrease in quality before transmission. In some embodiments, visual data can be video and or image data. In some embodiments, quality can relate to the resolution of visual data and/or other attributes such as a higher or lower frame rate.

In some embodiments, a combination of two or more example based models may provide an optimum resource for reconstruction of the lower-quality section of visual data. In some of these embodiments, the two or more example based models can be combined in such a way so as to form a new example based model that is transmitted alongside the lower-quality section of visual data. The lower-quality section of visual data can further be transmitted alongside more than one example based model in some embodiments, if that provides a required level of reconstruction into a higher-quality visual data.

In some embodiments, the live section of visual data transmitted will be at a lower quality. Aspects disclosed herewith can still be used in order to increase the quality of the section of visual data in some embodiments, even without an original high-quality version.

Optionally, one or more sample sections of visual data are transmitted from the first node to the second node in advance of a live transmission; and a subset of example based models is selected from the plurality of example based models based on a specified metric from the one or more sample sections of visual data; wherein the one or more example based

models used to substantially increase the quality of the live transmission is selected from the subset of example based models.

Optionally, one or more sample sections of visual data are transmitted from the first node to the second node in advance of a live transmission; wherein one or more example based models can be generated in advance of the live transmission based on a specified metric from the one or more sample sections of visual data.

Optionally, one or more sample sections of visual data are transmitted from the first node to the second node in advance of a live transmission; wherein one or more example based models from the plurality of example based models can be modified in advance of the live transmission based on a specified metric from the one or more sample sections of visual data.

The process of selecting the example based model to be used in some embodiments can be made more efficient by providing a smaller pool of example based models from which the selection can be made. Computational effort can thereby not be wasted in at least some of these embodiments, as a reconstruction using an example based model that is very unlikely to be used as the selected model for the relevant visual data, e.g. frame or video, will not be performed.

In some embodiments, by transmitting a section of visual data in advance of the live broadcast, example based models that are likely to be used can be grouped into a subset of the total plurality of example based models available.

During the live broadcast, in some embodiments, one or more example based models can be selected from the subset, rather than from the entire available library of models. If deemed necessary, in at least some embodiments, or if sufficient time is available before the live broadcast in other embodiments, the section of visual data can be transmitted in advance of the live broadcast and this can allow for the modification of an existing example based model or generation of a new example based model so as to provide a model or models for the more accurate reconstruction from a lower-quality section of visual data.

Optionally, the example based model comprises any of: a generative model; a non-linear hierarchical algorithm; or a convolutional neural network; or a recurrent neural network; or a deep belief network; or a dictionary learning algorithm; or a parameter; or a mapping function.

Optionally the example based model comprises a specific parameter or mapping function for the reconstruction of data.

In embodiments, the example based model can be one of a number of different model types. This increases the flexibility of the method and allows accurate reconstruction of the higher quality visual data in such embodiments.

Optionally the higher-quality visual data is divided into smaller sections based on similarities between a plurality of sections of the higher-quality visual data.

In some embodiments, by dividing the visual data into smaller sections or 'scenes', based on a particular metric for similarity, more efficient models can be generated. For example, in some embodiments multiple sections can be selected, all of which comprise part of a landscape shot, and only one model will be required to reconstruct the scene as opposed to a marginally different model for every separate section. If the next scene in the visual data is very different, for example a scene with significant movement, then in some embodiments this will be detected as such and a new model generated accordingly for the scene. Specific models can be trained for each scene or section in some embodi-

ments, enabling the reconstruction approach to be adaptive and to allow at least some optimisation of reconstruction.

Optionally a section of visual data can comprise any of: a single frame; or a sequence of frames; or a region within a frame or sequence of frames.

Depending on the visual data being processed, in some embodiments models can be generated for sections of visual data comprising a single frame, a sequence of frames or a region within a frame or sequence of frames. Each option could be necessary in some embodiments in order to provide the most efficient method of reconstructing a higher-quality version of a part of the original visual data.

Optionally the visual data is converted to a sequence of sections before the quality is reduced.

By separating the visual data into sections, for example by separating video data into a series of frames in some embodiments, the amount of data is temporarily increased, as the reductions in data from any compression methods previously used on the video file will be undone. However in some of these embodiments, this allows for the individual sections to be down-sampled thus reducing the size of the visual data, thereby allowing for lower-quality sections to be transmitted as re-encoded visual data in the original or optionally a more optimal codec but at a lower quality.

Optionally analysis of any two or more example based models occurs in parallel.

By analysing multiple example based models in parallel in some embodiments, the time taken to select a suitable example based model from a library of models can be reduced. This is especially beneficial for live broadcasting embodiments, where it can be advantageous for the time taken for processing visual data before transmission to be minimised.

In some embodiments, a pipeline of analysis can be performed for multiple subsequent sections of visual data in parallel. In some of these embodiments, analysis can be performed for subsequent sections of visual data in respect of compression settings and/or to classify the section or sections into categories in order to perform further content-specific analysis, where this further analysis can be performed once the visual data has been decoded, and in some embodiments this analysis can allow a reduction of the number of models considered for use with the visual data for enhancement or as starting points for training a model.

Optionally, a plurality of example based models are stored in a library.

In some embodiments, a stored library of example based models at the first node allows the analysis process to quickly select example based models for comparison without having to obtain them from the network. This, in turn, may result in a faster selection of the most suitable model for transmission across the network in some of these embodiments.

Optionally the sections of visual data are transmitted in the chronological order according to the section of higher-quality visual data.

By transmitting the sections of visual data in chronological order in some embodiments, the time difference between a section of visual data being recorded and a reconstructed higher-quality visual data being received by the second node can be minimised. For live broadcasting embodiments, this means that the lag between an event happening in front of a camera and a viewer watching a reconstructed version of the event can be reduced.

Optionally the sections of visual data are transmitted in the order in which they were transmitted by the first node.

Optionally the quality of the lower-quality sections of visual data is increased according to the order in which the sections of visual data were transmitted by the first node.

Some sections of visual data might be easier to process than others, for example a section of video comprising a single black frame would require less computational effort to process than a fast-moving scene. Therefore in some embodiments the sections of visual data that are easier to process could arrive at the second node before the more difficult to process sections. If the visual data is to be reconstructed and viewed in the order in which it was originally recorded in such embodiments, any sections of visual data which arrive out of sequence at the second node can be buffered in some of these embodiments and the sections placed in a chosen order, before increasing the quality and outputting the reconstructed visual data.

Optionally the analysis of an example based model for the use of reconstruction of the lower-quality section of visual data comprises the further steps of: extracting features of a section of the higher-quality visual data; and using the extracted features to select a predetermined example based model from the plurality of predetermined example based models to provide initial parameters for the developed example based model.

The closer the features of the example based model match those extracted from the corresponding section of visual data, the more accurate a reconstruction can be created using that particular example based model, in at least some embodiments. Therefore by extracting features from a section of the higher-quality visual data in some embodiments an appropriate model may be more accurately chosen from the plurality of example based models available.

Optionally the specified metric used to select and/or modify and/or generate an example based model is based on generating the highest quality output data following reconstruction, preferably wherein quality can be defined using any of: an error rate; a peak signal-to-noise ratio; or a structural similarity index.

A reconstructed section of visual data of the highest quality according to any of the above metrics can be desirable to a viewer, as it is liable to show the truest representation of the original higher-quality visual data, at least for some embodiments. In such embodiments, the broadcaster of the original higher-quality visual data is also likely to prefer that the viewer sees the visual data as it has been prepared, without distortion from the down sampling and transmission process.

Optionally the higher-quality section of visual data is divided into smaller sections based on similarities between a plurality of sections of the higher-quality visual data.

In some embodiments, by dividing the visual data file into smaller sections or 'scenes', based on a particular metric for similarity, less data needs to be transmitted over the network. For example, in some embodiments multiple sections can be selected, all of which comprise part of a landscape shot, and only one model will be required to reconstruct the scene to an acceptable level as opposed to a marginally different model for every separate section. If the next scene in the visual data is very different in such embodiments, for example a scene with significant movement, then it will be detected as such using the extracted features as described above and a new model can be selected accordingly.

Optionally a section comprises any of: a single frame, a sequence of frames and a region within a frame or a sequence of frames.

Depending on the visual data being processed, in some embodiments models can be selected for sections of visual

data comprising a single frame, a sequence of frames and a region within a frame or a sequence of frames. In these embodiments, each or all approaches could be necessary in order to provide the most efficient method of reconstructing a higher-quality version of a part of the original visual data.

Optionally the steps of transmitting the section of lower-quality visual data to the second node and transmitting to the second node the example based model that corresponds to the section of lower-quality visual data transmitted to the second node occur substantially simultaneously.

By transmitting the lower-quality section of visual data and the example based model simultaneously in some embodiments, the reconstruction process can be initiated substantially simultaneously on arrival of the data, as opposed to a delay between the arrival of the section of visual data and the example based model or vice versa. Both the example based model and the section of visual data are necessary for a higher-quality version of the section of visual data to be reconstructed, in some embodiments, and so any delay between their arrivals will result in a delay between the original visual data being transmitted from a first node and the reconstructed visual data being viewed at a second node.

Optionally the step of lowering the quality of a section of visual data is performed using down-sampling.

By lowering the quality of the section of visual data in some embodiments, less data needs to be sent across a network from a first node to a second node in order for a second node to receive the section of visual data. In these embodiments the lower-quality version, together with an example based model to be used for reconstruction, can still allow for less data to be transmitted than if original higher-quality section of visual data is transmitted.

Optionally a library is stored at both the first and second nodes, and a reference to one or more example based models from the library at the first node is transmitted to the second node.

It can advantageous in some embodiments to minimise the amount of data sent over a network. Higher-resolution videos can comprise significant amounts of data, which can be slow to transmit over existing network infrastructures such as the Internet. Some embodiments can provide a solution to the problem of video transmission over a network by transmitting instead a lower-quality section of video and an example based model, as described above. In some embodiments, in order to further reduce the amount of data transmitted over a network, a lower-quality section of visual data can be transmitted along with a reference to an example based model, in lieu of the example based model itself. In some of these embodiments the reference comprises less data than the example based model to which it refers. In such embodiments a library of similarly referenced example based models can be accessed from both the first and second nodes, so that once a model is selected at the first node and the reference transmitted, the correct model to be used at the second node can be identified.

Optionally, an off-site computing system performs a method comprising the steps of: receiving a lower-quality section of visual data; analysing a first example based model for the use of reconstruction of the lower-quality section of visual data; analysing a further one or more example based models for the use of reconstruction of the lower-quality section of visual data; selecting one or more example based models from a plurality of example based models to use based on a specified metric; generating a higher-quality section of visual data from the lower-quality section of

visual data using the one or more example based models that correspond to the lower-quality visual data.

Optionally, there is provided a method further comprising the step of: receiving a higher-quality section of visual data from a third node, before transmitting the higher-quality section of visual data to the first node.

Optionally, one or more of the nodes is based in an off-site computing system.

In some embodiments off-site, or 'cloud computing', systems allow for the performance of computerised tasks on a server not necessarily local to the site of the recording or reconstruction of a section of visual data. This allows for more powerful servers to be used according to the budget available for such services, and hence increased parallel processing of different example based models in some of these embodiments. The off-site system used in such embodiments could also provide a backup of any sections of visual data passing through the servers, thereby offering a solution in the case of loss of data at a site local to the recording or reconstruction of a section of visual data. If the computing system is scalable, as is preferable, then a growing amount of visual data processing work can be accommodated should the need arise in these embodiments.

Sub-Pixel Convolutions

According to an aspect, there is provided a method for enhancing one or more sections of lower-quality visual data using a hierarchical algorithm, the method comprising the steps of: receiving one or more sections of lower-quality visual data; enhancing the one or more sections of lower-quality visual data to one or more sections of higher-quality visual data using the hierarchical algorithm, wherein at least the first step of the hierarchical algorithm is performed in a lower-quality domain; and wherein the hierarchical algorithm operates in both a higher-quality domain and the lower-quality domain.

In embodiments, an algorithm performing at least the first processing steps of as enhancement in a lower quality domain ensures the visual data is at least partially if not substantially processed whilst still in lower-quality, reducing the computational complexity of the enhancement. Subsequent processing steps may move to the higher-quality domain to undertake further processing and output higher-quality visual data in such embodiments. Optionally, in some of these embodiments, only the last processing step or steps are performed in the higher-quality domain.

In some embodiments, resampling into the higher dimension space of visual data from a low-quality to high-quality domain happens before being processed through a super resolution network for enhancement, for example in some of these embodiments being enhanced using bicubic interpolation, but this enhancing does not add information useful for super resolution and forces the network to perform subsequent computation in a higher-quality domain when for example extracting feature maps and performing non-linear mapping. This "initial upscaling" makes the computation of the network/algorithm computationally expensive and increases memory requirements. In some embodiments, the low quality domain is a low resolution domain while the high quality domain is a high resolution domain.

Optionally, this method can be used as a filtering approach in place of other resampling or upsampling filters for use with encoding and decoding of visual data.

In some embodiments, using the method as a filter for visual data codecs can provide very high computational efficiency, and therefore can also provide minimal energy costs in performing such filtering. In these or other embodiments, the method can provide a filter that is fast and/or

flexible in expression and that can perform substantially accurate filtering in at least some embodiments.

Optionally, an activation function can follow the use of the method as a filtering approach.

Optionally, the first step of the hierarchical algorithm comprises extracting one or more features from the one or more sections of lower-quality visual data.

In embodiments, extracted features may be used to produce a value or series of values based on a metric from the input data. In some embodiments, the metric can then be used to select a hierarchical algorithm from the library which is most appropriate for the input data, as each hierarchical algorithm has associated metric values based on the input data from which the models were respectively trained, the selection based on the similarity between the metrics associated with the input data and each of the pre-trained models.

Optionally, the last step of the hierarchical algorithm comprises producing higher-quality visual data corresponding to the lower-quality visual data, and at least the last step of the hierarchical algorithm may be performed in the higher-quality domain.

In some embodiments, by using the last step of the hierarchical algorithm to perform enhancement from the lower-quality domain after feature maps have been extracted and non-linear mapping has been performed in the lower-quality domain, the layers other than the final layer are less computationally expensive and require less memory than if they were operating in high resolution space and the final layer can perform more complex upscaling.

In some embodiments, the final layer can be a sub-pixel convolution layer that is capable of upscaling low resolution feature maps into high resolution output, thus avoiding using a non-adaptive generic or handcrafted bicubic filter, and allowing for a plurality of more complex upscaling filters for each feature map. Further, in some of these embodiments, having the final layer perform super resolution from the lower resolution space can reduce the complexity of the super resolution operation.

Optionally, the higher-quality visual data is at a higher resolution than the lower-quality visual data, wherein the lower-quality visual data may contain a higher amount of artefacts than the higher-quality visual data.

In some embodiments, separating the visual data into a series of sections allows for the individual sections to be down-sampled thus reducing the visual data size, thereby allowing for lower quality sections to be transmitted as re-encoded visual data in the original or optionally a more optimal codec but at a lower resolution.

Optionally, the hierarchical algorithm comprises a plurality of connected layers, and the connected layers may be sequential.

Optionally, the at least one section lower-quality visual data comprises any of: a single frame of lower-quality visual data, a sequence of frames of lower-quality visual data, and a region within a frame or sequence of frames of lower-quality visual data. Furthermore, the lower-quality visual data may comprise a plurality of frames of video.

Depending on the visual data being processed, in some embodiment models can be generated for sections of visual data comprising a single frame, a sequence of frames or a region within a frame or sequence of frames. In these embodiments, some or each of these options could be used in order to provide the most efficient method of reconstructing a higher quality version of a part of the original visual data.

Optionally, the hierarchical algorithm differs for each section of visual data.

In some embodiments, each section of visual to be transferred is likely to be at least slightly different from the preceding section. These embodiments it can therefore allow for more accurate reconstruction if the hierarchical algorithm to be transmitted alongside each section of visual data is also different in order to account for these differences in the sections of visual data. For such embodiments, a new model can therefore be generated for each section of visual data accordingly.

Optionally, the hierarchical algorithm is selected from a library of algorithms.

In some embodiments, providing libraries of algorithms allows selection of an algorithm from the libraries and allows for modifications to be made to the stored hierarchical algorithms to enhance the visual data, reducing the computational complexity.

Optionally, the standardised features of the at least one section of received lower-quality visual data are extracted and used to select the hierarchical algorithm from the library of algorithms.

In some embodiments, extracted standardised features are used to produce a value or series of values based on a metric from the input data. In these embodiments, the metric can then be used to select the pre-trained model from the library which is most appropriate for the input data, as each model in the library has associated metric values based on the input data from which the models were respectively trained, the selection based on the similarity between the metrics associated with the input data and each of the pre-trained models.

Optionally, the hierarchical algorithm to be selected from the library of algorithms is based on generating the highest quality version of the lower-quality visual data, preferably wherein quality can be defined by any of: an error rate; a bit error rate; a peak signal-to-noise ratio; or a structural similarity index.

The predetermined metrics used in some embodiments to determine the hierarchical algorithm to be selected can be based on a predicted quality of the output data for each pre-trained model. In some of these embodiments, quality can be defined by any or all of: an error rate; a peak signal-to-noise ratio; or a structural similarity index.

Spatio-Temporal Approach

According to an aspect there is provided a method for enhancing at least a section of lower-quality visual data using a hierarchical algorithm, the method comprising the steps of: receiving at least a plurality of neighbouring sections of lower-quality visual data; selecting a plurality of input sections from the received plurality of neighbouring sections of lower quality visual data; extracting features from the plurality of input sections of lower-quality visual data; and enhancing a target section based on the extracted features from the plurality of input sections of lower-quality visual data.

In some embodiments, performing the enhancement on the extracted features reduces the computational expense of the enhancement by using hierarchical algorithms trained on the lower-quality visual data.

Optionally, the target section corresponds to one of the received plurality of neighbouring sections of lower-quality visual data. Alternatively, the target section may not correspond to one of the received plurality of neighbouring sections of lower-quality visual data.

In some embodiments, having the target section correspond to one of the received sections allows the enhancement to work in the spatial domain. Conversely, in some

embodiments, when the target section does not correspond to one of the received sections, then enhancement can be used to predict sections occurring between received sections, or alternatively to predict future sections.

Optionally, the at least one section of lower-quality visual data comprises any of: a single frame of lower-quality visual data, a sequence of frames of lower-quality visual data, and a region within a frame or sequence of frames of lower-quality visual data. The lower-quality visual data may further comprise a plurality of frames of video.

Depending on the visual data file being processed, in some embodiments models can be selected for sections of visual data comprising a single frame, a sequence of frames or a region within a frame or sequence of frames. In these embodiments each could be necessary in order to provide the most efficient method of reconstructing a higher quality version of a part of the original visual data.

Optionally, at least one of the plurality of selected input sections and the target sections are frames of lower-quality visual data, and the at least one of the plurality of selected input sections occurs sequentially before the target section. Furthermore, at least one of the plurality of selected input sections and the target sections may be frames of lower-quality visual data, and the at least one of the plurality of selected input sections may occur sequentially after the target section.

In some embodiments, having one of the selected input sections occur sequentially in time before or after the target section of the lower-quality visual data enables the enhancement to predict a possible future section, or section still to be received, but that is not included in the any possible target section. As such, in some of these embodiments, having a section before or after the target section may also enable the enhancement to predict a section occurring between two sections.

Optionally, the hierarchical algorithm comprises a plurality of layers, furthermore the layers may be any of sequential, recurrent, recursive, branching or merging.

Having a number of layers in some embodiments, which may or may not be sequential, recurrent, recursive, branching or merging allows different levels of processing to occur at different times and the layers can work in parallel, ensuring optimal efficiency when enhancing the lower-quality visual data.

Optionally, the extracting the subset of features is based upon a predetermined extraction metric.

In some embodiments, enabling a variety of features to be extracted as set by the predetermined extraction metric ensures the enhancement can be tailored to a particular system. As such, in some of these embodiments the computational expense can be increased or decreased appropriately based on system resources.

Optionally, the received plurality of neighbouring sections of lower-quality visual data are consecutive sections of lower-quality visual data.

In some embodiments the use of consecutive sections of lower-quality visual data enables more accurate enhancement of the lower-quality visual data. Receiving a plurality of sections or consecutive sections providing more information used to be able to be used when enhancing sections around the received sections of lower-quality visual data.

Optionally, a predetermined algorithm metric used to determine the hierarchical algorithm to be selected.

In some embodiments, having a plurality of predetermined algorithms which may be selected based on a predetermined algorithm metric ensure optimal efficiency when enhancing the visual data by having a library of hierarchical

algorithms which have already been determined to enhance the quality of visual data. In these embodiments, these predetermined algorithms may then be used as a basis for further enhancements.

Optionally, predetermined algorithm metric is based on generating the highest quality version of the lower-quality visual data, preferably wherein quality can be defined by any of: an error rate; a bit error rate; a peak signal-to-noise ratio; or a structural similarity index. Furthermore, standardised features of the received plurality of neighbouring sections of lower quality visual data may be extracted and used to select the hierarchical algorithm from the library of algorithms.

In some embodiments, a reconstructed section of visual data of the highest quality according to any of the above metrics can be desirable to a viewer, as it is liable to show the truest representation of the original higher-quality visual data. In these embodiments, the broadcaster of the original higher-quality visual data is also likely to prefer that the viewer sees the visual data as it has been prepared, without distortion from the down sampling and transmission process.

Optionally, other data can be provided to the algorithm. Optionally, the other data can be optical flow between frames. Optionally, the optical flow between frames can be calculated outside of the algorithm or can be part of the training of the algorithm.

In some embodiments, other data can be fed into the network, such as optical flow between the frames or a measure thereof. In these embodiments, the optical flow can be calculated separately outside the neural network or can be incorporated into the training of the neural network. In these embodiments, this data can simply be concatenated onto the end of each frame data. In some embodiments it is not sent during transmission. In other embodiments, a variant of optical flow can be extracted from the encoded video (as it forms a basis of the inter-frame compression).

Enhancing Visual Data

According to an aspect there is provided a method for enhancing at least a section of lower-quality visual data, the method comprising steps of: receiving at least a section of the lower-quality visual data; selecting a hierarchical algorithm from a plurality of hierarchical algorithms, wherein the step of selection is based on a predetermined metric and wherein the hierarchical algorithms were developed using a learned approach and at least one of the hierarchical algorithms is operable to increase the quality of the lower-quality visual data; and using the selected hierarchical algorithm to increase the quality of the lower-quality visual data to create a higher-quality visual data.

Some embodiments provide a method for enhancing at least a section of visual data that has been transmitted across a network. In these embodiments, by selecting a hierarchical algorithm based on a predetermined metric, a suitable hierarchical algorithm can be chosen from a library of algorithms that will be able to more accurately recreate the original visual data than a generic hierarchical algorithm.

Optionally, the lower-quality visual data is divided into smaller sections based on similarities between a plurality of sections of the original visual data.

In some embodiments, by dividing the visual data into smaller sections or 'scenes', based on a particular metric for similarity, more efficient models can be generated. For example, in some embodiments multiple sections can be selected, all of which comprise part of a landscape shot, and only one hierarchical algorithm will be required to remove artefacts from the scene as opposed to a marginally different model for every separate section. In some of these embodi-

ments, if the next scene in the visual data is very different, for example a scene with significant movement, then it will be detected as such and a new hierarchical algorithm can be generated accordingly. In some embodiments, specific hierarchical algorithms can be trained for each scene or section, enabling the artefact removal approach to be adaptive and to allow at least some optimisation of the artefact removal.

Optionally, the predetermined metric is based on a content-type and a measure of artefact severity of the lower-quality visual data. Furthermore, the artefact severity may be based on at least one artefact which may be any one of, or combination of: blocking, blurring, pixelation, ringing, aliasing, missing data, and other marks, blemishes, defects, and abnormalities in visual data.

In some embodiments, by basing the predetermined metric on the artefact severity of the lower-quality visual data, the artefact correction model resulting from the training and optimisation process can be referenced by the artefact severity of the scene it relates to. In these embodiments, this allows for a more accurate selection of a suitable artefact removal model when selecting a model from the library to clean new visual data.

Optionally, the plurality of hierarchical algorithms is any of: pre-trained hierarchical algorithms on example data; or manually created hierarchical algorithms.

In some embodiments, a pre-trained hierarchical or manually created hierarchical algorithm will be able to more accurately recreate the original visual data than a generic artefact removal hierarchical algorithm

Optionally, if a suitable hierarchical algorithm is not present in the plurality of hierarchical algorithms, a generic hierarchical algorithm is selected instead.

In some embodiments, a generic hierarchical algorithm may be used instead of a specific, trained model in situations where the library of algorithms does not include a visual data artefact removal model that has been trained on a similar enough content type and artefact graded to the received visual data. In these embodiments, the generic hierarchical algorithm may then be more accurate in removing the artefacts than the specifically trained hierarchical algorithms available.

Accelerating Machine Optimisation Processes

According to an aspect, there is provided a method for training learned hierarchical algorithms, the method comprising the steps of: receiving input data; generating metrics from the input data; selecting at least one hierarchical algorithm from a plurality of predetermined hierarchical algorithms based on comparing the generated metrics from the input data and like metrics for each of the plurality of predetermined hierarchical algorithms; developing the selected hierarchical algorithm based on the input data; and outputting a developed hierarchical algorithm.

Providing a plurality of pre-trained sets of parameters or models in some embodiments allows the training of a machine learning process to be accelerated. In these embodiments, input data to be processed by a trained machine learning process can have a tailored model, or set of parameters, developed for that input data based on the selected most similar pre-trained model. In such embodiments, the selection of the one or more most similar pre-trained model(s) can be done based on one or more metrics associated with the pre-trained models compared to the input data. In at least some of these embodiments the most similar pre-trained model can be used as a starting point for developing the tailored model for the data as a tailored model does not have to undergo as extensive development as needed when developing a model from first principles. In some

embodiments, using a similar or the most similar pre-trained model can also help with avoiding or escaping local minima (for incorrect solutions) during training.

Optionally, the method further comprises the step of storing the developed hierarchical algorithm with plurality of predetermined hierarchical algorithms along with the generated metrics from the input data.

In some embodiments, each new tailored model is then stored in a library of pre-trained models. In these embodiments, repeating or identical input data can re-use a tailored model that already exists without any further modification. In such embodiments, the closer a model is to an optimised version the less computation is required to generate a more optimal model. In these embodiments, as more signals and data are processed, an increasingly more extensive collection of models can be established in the library, requiring reduced development for tailored models for new data thereby saving computational effort.

Optionally, the plurality of predetermined hierarchical algorithms may be any of: pre-trained hierarchical algorithms on example data; or manually created hierarchical algorithms.

In some embodiments, the predetermined hierarchical algorithms can be based on example data or manually created to allow a user to create specific behaviours in any eventual trained model or set of parameters.

Optionally, the step of developing the selected hierarchical algorithm may be based on the input data comprises developing a more optimised hierarchical algorithm.

In some embodiments, developing the predetermined hierarchical algorithms can improve the example-based set of parameters to create a more optimised set of parameters.

Optionally, the hierarchical algorithm comprises at least one of any of: a specific parameter for the reconstruction of data; or a mapping function for the reconstruction of data.

In some embodiments, each hierarchical algorithm can be used to provide for the reconstruction of data. In some of these embodiments, the reconstruction process can vary, and therefore the model can comprise different features, coefficients and/or parameters that would allow for reconstruction and each layer of the hierarchical algorithm can comprise features, coefficients and/or parameters. Further, in some embodiments, the parameters for the reconstruction of visual data can comprise one or more hierarchical algorithms.

Optionally, the plurality of hierarchical algorithms may be stored in a repository.

For ease of access to the plurality of predetermined example-based models, in some embodiment the models can be stored in a repository. In these embodiments, the repository can take the form of a library. In some embodiments, the repository can be located on a remote device and/or distributed computing system (such as a cloud computing and/or storage service and/or distributed networked system or peer-to-peer system).

Optionally, the method may comprise the further steps of: extracting features of the input data; and using the extracted features to select a predetermined hierarchical algorithm from the plurality of predetermined hierarchical algorithms to provide initial parameters for the developed hierarchical algorithm.

In some embodiments, extracted standardised features are used to produce a value or series of values based on a metric from the input data. In these embodiments, the metric can then be used to select the pre-trained model from the library which is most appropriate for the input data, as each model in the library has associated metric values based on the input data from which the models were respectively trained, the

selection based on the similarity between the metrics associated with the input data and each of the pre-trained models.

Optionally, the metrics may be based on generating the highest quality output data, preferably wherein quality can be defined using any of: an error rate; a peak signal-to-noise ratio; a distance metric; a similarity measure; or a structural similarity index.

In some embodiments, the predetermined metrics used to determine the hierarchical algorithm to be selected is based on a predicted quality of the output data for each pre-trained model. In some of these embodiments, quality can be defined by any or all of: an error rate; a peak signal-to-noise ratio; or a structural similarity index.

Optionally, the hierarchical algorithm is selected based on providing optimal initialisation data when in use

By providing initialisation data which is optimal in at least some embodiments, thereby being as close as possible to the section of visual data being analysed, fewer iterations are required to adapt the model into a form to allow for the highest quality reconstruction according to any of the metrics above. In these embodiments, computational effort and time is therefore reduced when processing the section of visual data.

Enhancing Model Libraries

According to an aspect, there is provided a method for enhancing visual data when communicating visual data over a network from a first node to a second node, the method at the first node comprising the steps of: developing at least one modified hierarchical algorithm from a known hierarchical algorithm operable to substantially recreate at least one section of higher-quality visual data; transmitting to the second node references to one or more known hierarchical algorithms from which the modified hierarchical algorithms were developed; and transmitting to the second node one or more modifications to the one or more known hierarchical algorithms operable to reproduce the one or more modified hierarchical algorithms from the known hierarchical algorithms; wherein the second node is able to recreate substantially the higher-quality visual data using the modified hierarchical algorithm.

According to another aspect, there is provided a method for enhancing visual data when communicating visual data over a network from a first node to a second node, the method at the second node comprising the steps of: receiving a reference to a known algorithm corresponding to a section of higher-quality visual data; receiving corresponding modifications to the known hierarchical algorithm operable to produce a modified hierarchical algorithm, the modified hierarchical algorithm developed with the knowledge of the section of higher-quality visual data; and using the modified algorithm to substantially recreate the section of higher-quality visual data.

In some embodiments, knowledge of the section of higher-quality visual data can comprise knowing the encoder settings and/or the content category for at least one of the low-quality visual data, original visual data and the higher-quality visual data. In some of these embodiments, the content category can be based on a predetermined set of categories or a determines set of categories.

According to a further aspect, there is provided a system for enhancing visual data when communicating visual data over a network, the system comprising two or more nodes, wherein a first node is configured to: develop at least one modified hierarchical algorithm from a known hierarchical algorithm operable to substantially recreate at least one section of higher-quality visual data; transmit to the second node references to one or more known hierarchical algo-

gorithms from which the modified hierarchical algorithms were developed; and transmit to the second node one or more modifications to the one or more known hierarchical algorithms operable to reproduce the one or more modified hierarchical algorithms from the known hierarchical algorithms; wherein a second node is configured to: receive a reference to a known algorithm corresponding to a section of higher-quality visual data; receive corresponding modifications to the known hierarchical algorithm operable to produce a modified hierarchical algorithm, the modified hierarchical algorithm developed with the knowledge of the section of higher-quality visual data; and use the modified algorithm to substantially recreate the section of higher-quality visual data.

In some embodiments, by transmitting a reference to a known algorithm and modifications to that algorithm that allow a modified algorithm to be produced to enable reconstruction of a high-quality section or series of sections of visual data, less data can be transferred over a network to enable a high-quality visual data to be viewed compared to transmitting the high quality visual data. In some of these embodiments the algorithms, also known as models, representations or functions, can also be more accurate and therefore enable a more accurate reconstruction, i.e. produce a higher quality high-quality visual data, for example where quality can be measured by a low error rate in comparison to the original high-quality visual data.

Optionally, known hierarchical algorithm is selected from a library of hierarchical algorithms stored at the first node.

Optionally, the reference received by the second node corresponds to a known hierarchical algorithm stored in a substantially matching library at the second node.

Optionally, if the second node cannot identify the known hierarchical algorithm, the second node sends a request to the first node for transmission of the known hierarchical algorithm from the first node to the second node.

In some embodiments, by providing common or synchronised libraries of algorithms at both the first and the second nodes, and by transmitting a reference or reference identifier when transmitting the corresponding lower-quality visual data to allow selection of matching algorithms from both libraries using the reference identifier to identify the selected algorithm, only the reference identifier, the lower-quality visual data and the modifications to the stored hierarchical algorithms need to be transmitted between the nodes thus data transmission is reduced as the full algorithm doesn't need to be transmitted.

Optionally, the modified hierarchical algorithm is stored in a library at the first node after being developed.

In some embodiments, by storing the newly developed algorithm at the first node, the algorithm may be used as a starting point to develop further modified algorithms for similar high quality visual data, increasing the efficiency of the method in future situations.

Optionally, the modified hierarchical algorithm is stored in a library at the second node after transmission of the one or modifications to the known hierarchical algorithm from the first node to the second node.

In some embodiments, storing the modified algorithm at the second node after it has been recreated from the known algorithm and the transmitted modifications allows it to be referenced in the future when the modified algorithm is used as a starting point to develop a new algorithm at the first node. In such embodiments, this can prevent duplication.

In some embodiments, the first node may transmit known algorithms to the second node prior to sending the modifications required to make the modified algorithm to ensure

that the required known algorithm is present at the second node. In such embodiments, this can reduce errors.

Optionally, one or more representations of the one or more sections of higher-quality visual data are transmitted to the second node and to which the one or more modified hierarchical algorithms are applied to substantially reproduce the one or more sections of higher-quality visual data.

In some embodiments, by transmitting representations of the higher quality visual data to the second node the reconstruction of the higher quality visual data can be aided by applying the modified reconstruction model to the representation, but still with less data being transferred across the network than if the higher quality visual data had been transmitted.

Optionally, one or more sections of lower-quality visual data corresponding to the one or more sections of higher-quality visual data are transmitted to the second node and to which the one or more modified hierarchical algorithms are applied to substantially reproduce the one or more sections of higher-quality visual data.

Optionally, down-sampling is used to reduce the quality of one or more sections of higher-quality visual data to one or more sections of lower quality visual data.

Optionally, the one or more sections of lower-quality visual data are generated from the one or more sections of higher-quality visual data using a process of compression and/or quantisation.

In some embodiments, by lowering the quality of the section of visual data, less data needs to be sent across a network from a first node to a second node in order for the second node to receive the file. In these embodiments, the lower quality version together with a model to be used for reconstruction can still allow for less data to be transmitted than if the original higher-quality version of the same section of visual data is transmitted (where the quality can be determined by the resolution and/or the level of compression).

Strided Convolutions

According to an aspect, there is provided a method for enhancing at least a section of lower-quality visual data using a hierarchical algorithm, the method comprising steps of: receiving at least one section of lower-quality visual data; extracting a subset of features, from the at least one section of lower-quality visual data; forming a plurality of layers of reduced-dimension visual data from the extracted features; and enhancing the plurality of layers of reduced dimension visual data to form at least one section of higher-quality visual data wherein the at least one section of higher-quality visual data corresponds to the at least one section of lower-quality visual data received.

In some embodiments, extracting a subset of features from a section of lower quality visual data enables performance to be increased by only performing the enhancement on the subset of features as opposed to the entirety of the lower quality visual data, thereby reducing the amount of computation, whilst still achieving acceptable enhancements.

Optionally, the extracted subset of features is dependent upon one or more predetermined factors.

In some embodiments, the use of a predetermined factor to extract the subset of features enables the method to be tailored to a plurality of systems of different capabilities, and allows customisation of the amount of data to extract, allowing for a more accurate selection of a reconstruction model to enhance the lower quality visual data in an appropriate amount of time.

Optionally, the number of layers is dependent upon the one or more predetermined factors.

In some embodiments, the number of layers being based upon the one or more predetermined factors ensures the reconstructed higher-quality visual data corresponds to the lower-quality visual data received.

Optionally, the lower-quality visual data comprises a plurality of sections of video data.

In some embodiments, video may be used in the method, enabling the techniques to be applied to lower-resolution video data and upscaling it to corresponding higher resolution video data.

Optionally, the at least one section of lower-quality visual data comprises any of: a single frame of lower-quality video data, a sequence of frames of lower-quality video data, and a region within a frame or sequence of frames of lower-quality video data.

In some embodiments, depending on the visual data being processed, models can be used for sections of video data comprising a single frame, a sequence of frames or a region within a frame or sequence of frames. In some of these embodiments, each option could be necessary in order to provide the most efficient method of reconstructing a higher resolution version of a part of the received video file. In some embodiments, regions can include areas within a section, or segmentation between the background and foreground within a section, or some other content-based segmentation such as saliency based segmentation (i.e. where on each section of visual data a viewer is most likely to look).

Optionally, the hierarchical algorithm differs for each section of visual data.

In some embodiments, the use of different hierarchical algorithms for each section of visual data enables the most efficient hierarchical algorithm to be used for a particular section as opposed to using a single hierarchical algorithm for the entire visual data.

Optionally, the hierarchical algorithm is selected from a library of algorithms.

In some embodiments, a library of a plurality of algorithms allows for efficient selection of an hierarchical algorithm from the plurality of algorithms without delays due to communication links.

Optionally, standardised features of the at least one section of received lower-quality visual data are extracted and used to select the hierarchical algorithm from the library of algorithms.

Optionally, the hierarchical algorithm to be selected from the library of algorithms is based on generating the highest quality version of the lower-quality visual data, preferably wherein quality can be defined by any of: an error rate; a bit error rate; a peak signal-to-noise ratio; or a structural similarity index.

In embodiments, a variety of options are available in order to select which algorithm in the library of algorithms will output the best enhanced and enhanced visual data from the received visual data, for example an error rate such as the bit error rate can be used

General Aspects

It should be noted that in some aspects and/or embodiments, the terms model and/or algorithm and/or representation and/or parameters and/or functions can be used interchangeably.

It should also be noted that visual data, in some embodiments, may comprise image and/or video data.

Optionally, the modified hierarchical algorithm is trained using a learned approach.

In some embodiments, hierarchical or non-hierarchical algorithms can be substantially accurate and therefore enable a more accurate reconstruction, i.e. produce higher quality high-quality visual data from the low-quality visual data that is transmitted, for example where quality can be measured by resolution or by a low reproduction error rate in comparison to the original high-quality visual data. In some embodiments, using a learned approach can substantially tailor the hierarchical model or models for each portion of low-quality visual data.

Optionally, the learned approach comprises machine-learning techniques. The modified hierarchical algorithm may also be a non-linear hierarchical algorithm, which may comprise one or more convolutional neural networks.

In some embodiments, non-linear models can be substantially accurate in reconstructing higher-quality sections of visual data than dictionary-based approaches. In these embodiments, through using a learning-based approach, i.e. an approach that does not rely on pre-defined visual data features and operators, the model(s) can be optimised for each section or sequence of sections.

In some embodiments, the training of convolutional neural networks can be more computationally complex than dictionary learning for a similar accuracy, but the resulting model or algorithm can also be more flexible in representing visual data while using fewer coefficients for the reconstruction. In these embodiments, the resultant convolutional neural network model to be transmitted alongside the lower-quality visual data can be both smaller and can be more accurate in the reconstruction of a higher-quality visual data.

Some aspects can provide an improved technique for generating reconstruction parameters that can be used, when converting an original high-quality video into a down-sampled low-quality video, to allow recreation of a higher-quality version of the video from down-sampled low-quality video without significant loss in quality, for example having a low reconstruction error in comparison with the original high-quality video, and with a reduction in visual data transferred over a network. In such aspects, the application of such a technique can reduce the data transmitted when transmitting visual data in comparison with existing techniques while enabling reproduction of the visual data at its original quality without significant loss in quality in comparison to the original visual data (where quality can be defined by objective metrics such as error rate, PSNR and SSIM as well as subjective measures). In such aspects, such a proposed technique can allow minimal changes to be made to the overall infrastructure of video service providers, as it can augment most existing video compression techniques, and can provide advantages in video encoding and video streaming applications.

Optionally, the higher-quality visual data is at a higher-quality than the lower-quality visual data.

Optionally, the lower-quality visual data contains a higher amount of artefacts than the higher-quality visual data.

Optionally, the hierarchical algorithm performs visual data enhancement, preferably using super-quality techniques.

Optionally, the hierarchical algorithm uses a spatio-temporal approach.

In some embodiments, optionally for use for a section of visual data, the example based model may be a neural network and can use spatio-temporal convolution. In some embodiments, separating visual data into a series of sections allows for the individual sections to be down-sampled thus reducing the visual data size, thereby allowing for lower quality sections to be transmitted as re-encoded visual data

in the original or optionally a more optimal codec but at a lower quality. In some embodiments, a spatio-temporal network can allow an improvement in performance by exploiting the temporal information in the visual data and, for example, within a similar scene in sequential sections of visual data, there may be stationary sections of background in the sequential sections providing information relevant for the higher-quality version of that scene such that temporally consecutive sections can be used to super resolve one section.

Aspects and/or embodiments include a computer program product comprising software code to effect the method and/or apparatus of other aspects and/or embodiments herein described.

References to visual data can be references to video data and/or image data in some aspects and/or embodiments and vice versa. References to low-quality and/or lower-quality can be references to low-resolution and/or lower-resolution in some aspects and/or embodiments and vice versa. References to high-quality and/or higher-quality and/or highest quality and/or original quality can be references to high-resolution and/or higher-resolution and/or highest-resolution and/or original resolution in some aspects and/or embodiments and vice versa. References to sections can be references to frames and/or portions of frames in some aspects and/or embodiments and vice versa. References to enhance or enhancement can be references to upscale and/or upscaling in some aspects and/or embodiments and vice versa.

BRIEF DESCRIPTION OF DRAWINGS

Embodiments will now be described, by way of example only and with reference to the accompanying drawings having like-reference numerals, in which:

FIG. 1 illustrates an over-complete dictionary of 16×16 atoms;

FIG. 2a illustrates the layers in a convolutional neural network with no sparsity constraints;

FIG. 2b illustrates the layers in a convolutional neural network with sparsity constraints;

FIG. 3 illustrates the encoding process to generate for transmission, from high-resolution visual data, a combination of low-resolution visual data and convolutional neural network able to use super resolution to increase the resolution of the low-resolution data;

FIG. 4 illustrates the decoding process using the low-resolution visual data and convolutional neural network to recreate a version of the high-resolution visual data;

FIG. 5 is a flowchart illustrating the method steps for encoding a visual data for transmission using the process of FIG. 3;

FIG. 6 is a flowchart illustrating the method steps for decoding the visual data and convolutional neural network data generated using the method shown in FIG. 5;

FIG. 7 is a diagram of an efficient sub-pixel convolutional network according to an embodiment, featuring two feature map extraction layers built with convolutional neural networks, and a sub-pixel convolution layer that aggregates the feature maps from low-resolution space and builds the super resolution image in a single step;

FIG. 8 is a diagram of the network architecture according to an embodiment for super resolution using three input frames where only the input layer is modified to jointly handle three input frames for the super resolution prediction of only the middle frame while the hidden layers are identical to those in FIG. 7;

FIG. 9 is a flow chart illustrating the method steps for a machine learning process;

FIG. 10 illustrates the encoding process to generate for transmission, from high resolution visual data and known convolutional neural networks, low resolution visual data and modifications to the known convolutional neural network;

FIG. 11 illustrates the decoding process using the low resolution visual data and a reference to a known convolutional neural network, along with modifications to that network, to recreate a version of a high resolution visual data;

FIG. 12 illustrates a further encoding process to generate for transmission, from high resolution visual data and known convolutional neural networks a modified convolutional neural network; and

FIG. 13 illustrates a further decoding process using a reference to a known convolutional neural network, along with modifications to that network, to recreate a version of a high resolution visual data;

FIG. 14 is a flowchart illustrating the method steps for a conventional image enhancement process from a lower resolution original visual data to a higher resolution visual data;

FIG. 15 is a flowchart illustrating the method steps for an image enhancement process according to an embodiment, where visual data is up-scaled from a lower-resolution original visual data to a higher-resolution visual data, using a combination of the received visual data and a convolutional neural network selected from a library, the convolutional neural network able to use super resolution techniques to increase the resolution of the received visual data;

FIG. 16 is a flowchart illustrating the method steps for an image enhancement process to: generate a down-sampled visual data for transmission from a higher-resolution original visual data; the transmission of the down-sampled visual; and an image enhancement process to upscale from the down-sampled visual data to a higher-resolution visual data, using a combination of the received down-sampled visual data and a convolutional neural network selected from a library, the convolutional neural network able to use super resolution techniques to increase the resolution of the received down-sampled visual data;

FIG. 17 is a flowchart illustrating the method steps for an image enhancement process according to the embodiment in FIG. 16, with additional steps for encoding visual data at another node prior to transmitting the encoded visual data to the node performing the method of FIG. 16;

FIG. 18 is a flowchart illustrating the method of FIG. 15, with the additional steps of extracting a reduced dimension representation and concatenating the visual data to create a reduced dimension representation of the separated frame;

FIG. 19 is a flowchart illustrating the method of FIG. 16, with the additional steps of extracting a reduced dimension representation and concatenating the visual data to create a reduced dimension representation of the separated scene;

FIG. 20 is a flowchart illustrating the method steps for an image enhancement process according to the embodiment in FIG. 19, with additional steps for encoding visual data at another node prior to transmitting the encoded visual data to the node performing the method of FIG. 19;

FIG. 21 illustrates the selection of pixels from visual data based upon a predetermined value to produce reduced resolution visual data for processing by the method illustrated in FIGS. 18 through 20;

FIG. 22 illustrates an overview of a method of generating models for use in image artefact removal;

FIG. 23 illustrates a method of using the image artefact removal models;

FIG. 24 illustrates an alternative method of using the image artefact removal models;

FIG. 25 illustrates the encoding process to generate, for transmission, from high resolution video data, a combination of low resolution visual data and a reference to an algorithm able to use super resolution to increase the resolution of the low resolution visual data;

FIG. 26 illustrates the decoding process using the low resolution visual data and the algorithm corresponding to the received reference to recreate a version of the high resolution visual data; and

FIG. 27 illustrates the method steps of the encoding process to generate for transmission, from original visual data being recorded through to the transmitted package of a lower-resolution scene along with a hierarchical algorithm.

DETAILED DESCRIPTION

Various embodiments will now be described in detail with reference to the Figures referenced above.

Reduction of Complexity in Neural Networks

With reference to FIGS. 2a and 2b, various possible configurations of neural network for use in at least some embodiments shall now be described in detail.

An example layered neural network is shown in FIG. 2a having three layers 10, 20, 30, each layer 10, 20, 30 formed of a plurality of neurons 25, but where no sparsity constraints have been applied so all neurons 25 in each layer 10, 20, 30 are networked to all neurons 25 in any neighbouring layers 10, 20, 30. The example simple neural network shown in FIG. 2a is not computationally complex due to the small number of neurons 25 and layers. Due to the density of connections, however, the arrangement of the neural network shown in FIG. 2a won't scale up to larger sizes of network, i.e. the connections between neurons/layers, easily as the computational complexity soon becomes too great as the size of the network scales and scales in a non-linear fashion.

Where neural networks need to be scaled up to work on inputs with a high number of dimensions, it can therefore become too computationally complex for all neurons 25 in each layer 10, 20, 30 to be networked to all neurons 25 in the one or more neighbouring layers 10, 20, 30. A predetermined initial sparsity condition is used to lower the computational complexity of the neural network, for example when the neural network is functioning as an optimisation process, by limiting the number of connection between neurons and/or layers thus enabling a neural network approach to work with high dimensional data such as images.

An example of a neural network is shown in FIG. 2b with sparsity constraints, according to at least one embodiment. The neural network shown in FIG. 2b is arranged so that each neuron 25 is connected only to a small number of neurons 25 in the neighbouring layers 40, 50, 60 thus creating a neural network that is not fully connected and which can scale to function with, higher dimensional data—for example as an optimisation process for video. The smaller number of connections in comparison with a fully networked neural network allows for the number of connections between neurons to scale in a substantially linear fashion.

Alternatively, in some embodiments neural networks can be use that are fully connected or not fully connected but in different specific configurations to that described in relation to FIG. 2b.

Further, in some embodiments, convolutional neural networks are used, which are neural networks that are not fully connected and therefore have less complexity than fully connected neural networks. Convolutional neural networks can also make use of pooling or max-pooling to reduce the dimensionality (and hence complexity) of the data that flows through the neural network and thus this can reduce the level of computation required. In some embodiments, various approaches to reduce the computational complexity of convolutional neural networks can be used such as the winograd algorithm or low-rank matrix approximations.

End-to-End Super Resolution Enhancement Model Generation & Reconstruction

Although initially designed to enhance single image frames, super resolution techniques can also be used on multiple frames in some embodiments. To apply these techniques in such embodiments, multiple low-resolution frames can be gathered together and the sub-pixel shifts between the individual frames can be used to create a higher resolution image than the original. In such embodiments, a series of frames can be combined to form a higher resolution video than was originally provided.

To improve on the above described dictionary-based approach for encoding and decoding video using super resolution techniques, in at least some embodiments it is proposed to use deep learning techniques and convolutional neural network models instead of dictionary learning techniques and dictionary-based models.

Dictionary learning is closely related to sparse coding techniques, while deep learning is only more loosely connected with sparse coding. Sparse coding is an effective mechanism that assumes any natural image can be sparsely represented in a transform domain. Sparse coding is a technique used to automatically find a small number of representative patterns which, when combined in the right proportions, reproduce the original input patterns. The sparse coding for an input then consists of those representative patterns. From a signal processing perspective, it follows that complicated signals can be broken down into a small, i.e. sparse, number of simple signals.

The transform domain can be a dictionary of image atoms, which can be learnt through a training process known as dictionary learning that tries to discover the correspondence between low-resolution and high-resolution sections of images (or “patches”). Dictionary learning uses a set of linear representations to represent an image and where an over-complete dictionary is used, a plurality of linear representations can be used to represent each image patch to increase the accuracy of the representation.

In at least some embodiments, instead it is proposed to use machine learning with deep learning techniques that can instead create non-linear representations of an image or sequence of images.

When using machine learning and sparse coding principles, a training process is used to find optimal representations that can best represent a given signal, subject to predetermined initial conditions such as a level of sparsity.

When using convolutional neural networks in at least some embodiments, the efficiency in terms of computational and memory cost can be important.

In some embodiments, convolutional neural network models (or hierarchical algorithms) can be transmitted along with the low-resolution frames of video data because the convolutional neural network models reduce the data transmitted in comparison with learned dictionaries being transmitted along with the low-resolution images, especially in comparison to transmitting learned over-complete diction-

aries with the low-resolution images, in the case where the models and dictionaries have the same level of reconstruction accuracy.

The convolutional neural network models of some embodiments allow the reconstruction process and image representation to be non-linear. This is in contrast to the dictionary-based approach, which represents images in a linear fashion, as described above. The convolutional neural network models of some embodiments offer reconstruction with a feed-forward mechanism based on image convolution that is able to be computed in parallel and for which methods with high computational speed are available, such as the AMD® Tiled Convolution approach which can be used for fast image filtering.

The series of convolutions in a convolutional neural network model of some embodiments allow the neural network to be used to search for pixel correlations in a region far larger than the initial patch decomposition, but applying a weighting to weaken the importance of correlations for locations further away from the pixel of interest. In contrast, linear sparse coding approaches like dictionary learning are restricted to looking for correlations in an initial patch size, to avoid the computational complexity of searching the entire image. As a result, the method of these embodiments can more fully exploit the natural or non-local redundancies of the video frames and between a sequence of video frames.

In some embodiments, optionally the convolutional neural network models only assume local spatial or spatio-temporal correlations in an image or video. This assumption is in contrast to the assumption that small data patches must be well represented by a restricted number of atoms from a dictionary, which is a sparsity constraint applied to dictionary learning and over-complete dictionary approaches.

In some embodiments, recurrent neural network models can be used. Recurrent neural network models have single layers that are used iteratively. The layers of a recurrent neural network model can be unrolled into a number of layers in order to be treated in a similar fashion to convolutional neural networks, in some embodiments.

Referring now to FIGS. 3 and 5, a technique that can be used in method, apparatus and/or system embodiments to encode video data will now be described in detail.

Original video data **70** is a high-resolution video, for example having a resolution of 1920 pixels by 1080 pixels (also known as “1080p” video) or 3840 pixels by 2160 pixels (also known as “4K” video). This video data can be encoded in a variety of known video codecs, such as H.264 or VP8, but can be any video data that can be decoded into the component frames of the video. The original video data **70** is provided as input visual data for encoding using the technique.

The original video data **70** is then split into single full-resolution frames at step **80** (or step **190**), i.e. into a sequence of images at the full resolution and/or quality of the original video data **70**. For some video codecs, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

Optionally, at step **90** (or step **190**), in some embodiments the full-resolution frames can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection”. The video data is split or clustered into scenes to enable more specific training and optimisation. By scene, it is meant a consecutive group or sequence of frames, which at the coarsest level can be the

entire video, a single frame or at the most granular level can be or a section/segment of a frame.

The exact approach employed for splitting frames into scenes can be performed in many ways, depending on embodiment, with the four main approaches being described below:

A first approach is the use of the frame-type groupings of the underlying video codec. In this approach, a scene is defined in the encoded video by each group of pictures (sometimes referred to by the abbreviation “GoP”), for example between two I-frames. This is probably the simplest approach for scene selection and does not require much additional computation to perform scene analysis of the video, but the approach does not allow any flexibility in the level of granularity in the definition of each scene.

The second approach is the use of a scene or shot transition detector, for example as described in the paper “A Unified Model for Techniques on Video-Shot Transition Detection” by Jesús Bescós, Guillermo Cisneros, José M. Martínez, José M. Menéndez, and Julián Cabrera as published in IEEE TRANSACTIONS ON MULTIMEDIA, VOL. 7, NO. 2, April 2005 on pages 293-307 which is incorporated herein by reference. This approach proposes a unified detection model to detect camera-shot transitions in video, detecting both abrupt and gradual camera-shot transitions.

The third approach is the use of a clustering approach after applying an unsupervised learning or dimensionality reduction approach such as K-means clustering or manifold learning or Principle Component Analysis (PCA) etc. A suitable K-means clustering approach is detailed in the paper “Constrained K-means Clustering with Background Knowledge” by Kiri Wagstaff, Claire Cardie, Seth Rogers and Stefan Schroedl as published in the Proceedings of the Eighteenth International Conference on Machine Learning, 2001, p. 577-584. A suitable manifold learning approach is proposed in the paper “Algorithms for manifold learning” by Lawrence Cayton published on Jun. 15, 2005. Both of these papers are herein incorporated by reference.

The fourth approach is to use a range of predefined scene classes and then automatically classifying scenes using the predefined classes. Examples of this approach can be found in the paper “Nonparametric Scene Parsing: Label Transfer via Dense Scene Alignment” by Ce Liu, Jenny Yuen and Antonio Torralba having IEEE reference 978-1-4244-3991-1/09 and the paper “80 million tiny images: a large dataset for non-parametric object and scene recognition” by Antonio Torralba, Rob Fergus and William T. Freeman. Both of these papers are herein incorporated by reference.

The exact approach for how a scene is defined is independent from the rest of the approach, and will vary by embodiment, although it will have an impact on the training time and the reconstruction performance so there is a level of optimisation required for each possible embodiment to find a balance between the speed and the accuracy of the scene detection step.

Alternatively, no scene matching needs to be performed at step 90 (or step 190) in some embodiments. A further alternative for some embodiments is to perform only very basic time-stamping of the video, to determine how to reassemble the video frames for later stages in the technique, for example during transmission or during playback of the video.

Regardless whether the video has been broken into frames or scenes (i.e. groups of multiple frames) in step 90 (or step 190) or remains as a sequence of frames from step 80 (or step 190), each frame is down-sampled into lower resolution

frames at a suitably lower resolution. Optionally, in some embodiments this step can occur before the frames are grouped into scenes in step 80 (or step 190), so step 90 (or step 190) can be exchanged with step 90 (or step 190) in these alternative embodiments. The lower-resolution frame can, for example, be 33% to 50% of the data size relative to the data size of the original-resolution frame, but in should be appreciated that in embodiments the lower resolution can be any resolution that is lower than the original resolution of the video.

In step 110 (or step 200), in at least one embodiment super resolution techniques are employed to create a specific model using machine learning for each frame, so that the model can be used to substantially recreate the original resolution version of a lower-resolution frame and trained using machine learning based on knowledge of the original-resolution frame. This step is termed the training and optimisation process. In some embodiments, generic models can be developed for types of scene or frame. Alternatively, in other embodiments, models can be developed for each scene.

By employing a deep learning approach to generating the model in embodiments, a non-linear hierarchical model can be created in some of these embodiments to reconstruct a higher-resolution frame from the lower-resolution frame.

An example of a deep learning approach, but in respect of only still images and without use of the original image to optimise the reconstruction model, is described in the paper “Learning a Deep Convolutional Network for Image Super-Resolution” by Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang published in D. Fleet et al. (Eds.): ECCV 2014, Part IV, LNCS 8692, pp. 184-199, 2014 and this paper is incorporated herein by reference. This paper relates to using super resolution techniques when trying to obtain an unknown high resolution version of a given low resolution image and proposes the use of a convolutional neural network to learn mappings between low resolution and high resolution images.

While the creation of non-linear hierarchical models is more computationally complex than creating linear models using dictionary learning approaches, in at least some embodiments it can be more flexible. Further, the non-linear models created in some of these embodiments can use only a few coefficients to enable reconstruction in comparison to the dictionary-based approach, thus such models can require less data to be transmitted. Still further, the non-linear models used in some embodiments can be more accurate in reconstructing higher-resolution frames than dictionary-based approaches. The non-linear models created in some embodiments are small convolutional neural networks rather than over-complete dictionaries. In contrast to the local patch averaging approach that tends to be used in reconstruction by dictionary learning approaches, the use of a convolutional neural network model in embodiments also allows a more appropriate filter to be learned for final reconstruction where neighbouring patches are considered, which can avoid unintended smoothing of the reconstructed higher-resolution image.

The training and optimisation process in embodiments can be configured according to a desired trade-off between computational time spent and desired quality of results. In general, the number of iterations used during the training process yields approximately logarithmic gains in reconstruction accuracy, so in some embodiments it is preferred to use an automatic threshold to stop further optimisation. When favouring quality of results in embodiments, the automatic threshold can be set to a predetermined value of

reconstruction error, for example by calculating the mean squared error, but other methods can be used. Alternatively, in some embodiments the automatic threshold can be set to limit the training and optimisation process to a predetermined number of iterations. As a further alternative, in some 5 embodiments a combination of these two factors can be used.

In some embodiments, step **120** (or step **210**), the portion of the low-resolution video and the reconstruction model for that portion of video are output for transmission. Optionally, 10 in some embodiments the video and model can be stored together within a suitable data container format such as a Matroska Multimedia Container (otherwise known as a MKV container). Alternatively, in some embodiments the video and model can be combined with other sections or the 15 entire video and placed into a suitable data container format. In some embodiments, at step **120** (or step **210**), the low-resolution video frames can be re-encoded using either the original video codec applied to the original video data **70** or, alternatively, a more optimal video codec can be applied to 20 the video data to produce output video data **130**. Optionally, if scene detection or time stamping was performed, in some embodiments the data output for transmission can include either a list of scenes or time stamp data respectively, or this data could be stored within the data container.

In embodiments there can be a number of variations to the encoding framework described above. For the above described technique, it can be assumed in many embodi- 25 ments that there would be a one-to-one mapping between scene and model, however this does not need to hold true as a local cluster of scenes could be used to train each model in some embodiments. Conversely, it could also be possible in some embodiments to use several similar models to initialise (by way of a weighted average) initial parameters 30 for training new unseen scenes. It may also be sufficient to simply weight the models without requiring any further optimisation in some embodiments—similar to label fusion methods used in medical imaging for example.

Referring now to FIGS. **4** and **6**, embodiments for reconstructing the video encoded using the technique will now be 40 described in detail.

First, the data **130** is received from the network. In 45 embodiments, the data received **130** depends on how the data was prepared for transmission in step **120** (or **210**) as detailed above in the various possible embodiments and the prepared data can include the video data and one or more reconstruction models for that video data. In some embodi- 50 ments, it is possible that the video data and one or more reconstruction models are not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required to decode or reconstruct the higher-resolution video from the data transmitted over the network.

At steps **140** and **150** (and step **220**), the received data is prepared for reconstruction depending on the embodiment. 55 In most embodiments, this step involves separating the low-resolution video from the reconstruction models (step **220**). Optionally, in some embodiments the low-resolution video is decompressed using the video codec used for transmission into full-resolution image frames (step **230**) 60 and each of the frames is matched up with the corresponding reconstruction model, which are also unpacked into frame-by-frame order (step **240**).

At step **160** (step **250**), the reconstruction model is applied to each of the frames to output higher-resolution frames 65 depending on the embodiment. The reconstruction, or decoding, process in most embodiments involves applying

the optimised super resolution convolutional neural network model, or reconstruction model, for each scene in order to restore the lower-resolution video to its original resolution having substantially the same quality as the original high- 5 resolution video. Given the corresponding models for each lower-resolution frame, the original higher-resolution frames can be reconstructed with high accuracy for at least some of these embodiments.

The complexity of the reconstruction process can be far 10 simpler, i.e. require less computation, when compared to that of training the model in most embodiments. In most embodiments, the reconstruction process is very similar to a single iteration of the training process, however there is no optimisation required and the process simply applies the learnt 15 convolutional neural network model to reconstruct the image on a patch-by-patch basis. In some embodiments, a singular model is applied to each scene and can be sent as accompanying data in the video container ahead of the scene in which it is used. In some embodiments, the model 20 associated with each scene can either be indexed in a data stream, synchronised with the video and audio stream, or indexed in a frame-to-model look-up table which can be sent ahead of the video stream depending on the embodiment.

The image reconstruction in most embodiments has a 25 linear computational complexity, dependent on the total number of pixels of each frame, as the reconstruction process requires a series of convolution functions (typically less than 3) to be applied for each relevant image patch centred around each pixel. A basic reconstruction step for 30 each image is as described in the paper “Learning a Deep Convolutional Network for Image Super-Resolution” by Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang published in D. Fleet et al. (Eds.): ECCV 2014, Part IV, LNCS 8692, pp. 184-199, 2014, and this step can be used 35 with or without modification in some embodiments however, in some of these embodiments these functions can be applied in parallel across multiple images and in some embodiments it can be feasible to perform real-time reconstruction for video playback.

Super resolution operates on a patch-by-patch basis in 40 most embodiments, but it may not be necessary to perform super resolution for every single patch of an image in some embodiments. In some embodiments, where the patch is flat and smooth, using a bi-cubic upscaling can be sufficient. It may be faster in some embodiments, for both training and 45 reconstruction, if these patches could be bypassed. To identify these patches, one approach as set out in the paper “Coupled Dictionary Training for Image Super-Resolution” by Jianchao Yang, Zhaowen Wang, Zhe Lin, Scott Cohen and Thomas Huang (as published in IEEE TRANSACTIONS ON IMAGE PROCESSING, VOL. 21, NO. 8, AUGUST 2012), incorporated by reference, can be used in 50 some embodiments to filter out patches which have low variances and ignoring these areas of the image. Alternatively, in other embodiments edge detectors could be used instead to identify areas of interest. 55

Further still, in some embodiments some level of inter- 60 frame comparison can be used to identify regions of the image which do not change from frame to frame, thereby eliminating areas of the image from duplicating the efforts of super resolution reconstruction which occur in a previous frame. In these embodiments the underlying video code for the low-resolution video can be used to identify the regions of change as part of its inter-frame compression mechanism 65 and this information could be taken advantage for more efficient real-time super resolution reconstruction and video playback. In addition, in some embodiments an auxiliary

data-structure, such as a random fern, can be used to keep track of the uniqueness of patches that have been reconstructed. Given that an image may contain many repeating patches, use of this approach in some embodiments could prevent unnecessary reconstruction of patches whereby the result is already known (in another patch). Similar approaches have been used in key frame detection for real-time camera localisation for augmented reality applications.

The proposed method and system of encoding and decoding video of at least some of the embodiments can augment existing video compression and decompression codecs. By using the same codec as used to create the video data, the low resolution video that is transmitted can be encoded with the same video codec but at a lower resolution than that of the original in some embodiments. Thus, in some embodiments the proposed method and system can be independent of the original codec used to create and compress the video. Alternatively, in some embodiments it is possible to select an alternative codec for transmission and to re-encode the video back into the original codec when the video data has been transmitted, uncompressed and restored to a high resolution version of the original video data.

By incorporating, for example in some embodiments, image enhancement techniques into a video compression process, a reduction in resolution of the video during transmission can be achieved and thus lower amounts of data are transmitted because the image can be enhanced at the node receiving the reduced-resolution video to recreate the video at the original resolution. By using super resolution as the image enhancement technique in such embodiments, and with the knowledge of the original high-resolution image before applying image enhancement techniques, it is possible to down-sample the original high-resolution image to a lower-resolution image and choose one or more representations from a library of representations that can be used to create a high-resolution image from the down-sampled image based on the knowledge of the original image. In such embodiments, this results in only a down-sampled image and the super resolution reconstruction model or parameter being transmitted over a network instead of the original high resolution video, thus decreasing the amount of data transferred while maintaining a similar level of image quality relative to the original once the high-resolution image has been recreated.

At step **170**, in some embodiments each of the segments of video are output from the reconstruction process as higher-resolution frames at the same resolution as the original video **70**. In these embodiments, the quality of the output video **180** is substantially similar to that of the original video **70**, within the error bounds optionally applied by the machine learning process that develops the reconstruction model at step **110** (or step **200**) of the encoding process.

At step **180**, in some embodiments the segments of video are combined such that the video can be displayed. Optionally at step **170** or **180**, in other embodiments the video can be re-encoded with the original codec used on the original video **70** or a more predetermined optimal codec, to enable playback on a decoder or display device.

In most embodiments, the proposed technique does not perform conventional compression by rearranging the original data and discarding redundant data or data deemed to be of least importance such that the majority of the data is intact. Instead, in these embodiments the technique determines a series of simple light-weight functions through machine learning that can be used to reconstruct the original data while minimising the reconstruction error given some

pre-defined cost function (a number of optimisation or search frameworks common to machine learning can be used, such as gradient descent or stochastic gradient descent). Thus, it follows that the technique of these embodiments can be described as an optimisation process of the reconstruction functions rather than the data redundancy/reduction process of conventional compression techniques. Sub-Pixel Convolution

Referring now to FIGS. **3** and **5**, embodiments using another technique to encode visual data will now be described in detail. These embodiments can be used in combination with other embodiments and alternative and optional embodiments described elsewhere in this specification.

Again, in these embodiments, original video data **70** is provided into the method or system and is a high-resolution video, for example having a resolution of 1920 pixels by 1080 pixels (also known as “1080p” video) or 3840 pixels by 2160 pixels (also known as “4K” video). This video data can be encoded in a variety of known video codecs, such as H.264 or VP8 but can be any visual data for which the system or method is able to decode into the component sections.

The original video data **70** is then split into single full-resolution frames at step **80** (or step **190**), i.e. into a sequence of images at the full resolution of the original video data **70**. For some video codecs, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

Optionally, at step **90** (or step **190**), the full-resolution frames can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection”. The video data is split or clustered into scenes to enable more specific training and optimisation. By scene, it is meant a consecutive group or sequence of frames, which at the coarsest level can be the entire video or at the most granular level can be a single frame.

To super resolve a low-resolution image to high-resolution space, it is necessary to increase the dimensionality of the low-resolution image to match that of the high-resolution image at some point during the operation of the example based model or convolutional neural network. By upscaling the low-resolution image to a high-resolution as an input to an example-based model, the example-based model does not need to learn the super resolution filters to perform upscaling from the low-resolution space to a high-resolution space but this approach can sacrifice reconstruction accuracy and introduces computational complexity as all subsequent processing must occur at the high resolution dimensionality.

Referring to FIG. **7**, there is shown an efficient sub-pixel convolutional neural network (ESPCN) **700** having a low-resolution input image **710** with two feature map extraction layers **720**, **730** built with convolutional neural networks and a sub-pixel convolution layer **740** that aggregates the feature maps from low-resolution space and builds the super resolution image **750** in a single step.

In an embodiment, as shown in FIG. **7**, the high-resolution data is super resolved from the low-resolution feature maps only at the very end of the network, in the final or last layer of the network, therefore avoiding the need to perform most of the computation in the larger high-resolution dimensionality. As processing speed depends substantially directly on the input image dimensionality, operation in the low-resolution space will allow for faster processing. Further, by performing feature map extraction before upscaling, more complex and more accurate upscaling (i.e. additional recon-

struction accuracy) is possible by the final or last layer of the network as an upscaling or resampling function can be learned and is content and data adaptive as compared to the applying an explicit generic or handcrafted interpolation filter to perform upscaling (for example, a bicubic interpolation filter) that is heuristically chosen beforehand.

The task of the example based model, in some embodiments a single image super resolution network, is to estimate a high resolution image given a low resolution image that is downsampled or downsampled from a corresponding high resolution image. In some cases, for example video compression, the down sampling operation can be deterministic and known: to produce the low resolution image from the high resolution image, the high resolution image can be convolved using a Gaussian filter (but other filters can be used), thus simulating a point spread function in a camera, then downsampled by a upscaling ratio r . However general super resolution and upscaling are typically not deterministic and in some embodiments, in order to mitigate for non-deterministic operations, training can be performed for a range of different down sampling operations. In other embodiments, estimation can be performed, for example by training a classifier or some visual comparison, and optionally this could be parametrised for the super resolution process in some embodiments. In general, both the low and high resolution image have C colour channels, thus can be represented as real-valued tensors of size $H \times W \times C$ and $rH \times rW \times C$ respectively. In some embodiments, rather than recovering an high resolution image from an upsampled and interpolated version of the low resolution image, instead the low resolution image is used by a three layer convolutional network to avoid upscaling before the low-resolution image is fed into the network. In the network, a sub-pixel convolution layer is applied to upscale the low resolution feature maps to produce a high resolution image using super resolution.

For a network composed of L layers, the first $L-1$ layers can be described as follows:

$$f^l(I^{LR}; W_l, b_l) = \mathcal{O}(W_l * I^{LR} + b_l)$$

$$f^l(I^{LR}; W_{l,b}, b_{l,i}) = \mathcal{O}(W_l * f^{l-1}(I^{LR}) + b_l)$$

where $W_l, b_l, l \in (1, L-1)$ are learnable network weights and biases respectively. W_l is a 2D convolution tensor of size $n_{l-1} \times n_l \times k_L \times k_L$, where n_l is the number of features at level l , $n_0 = C$, and k_L is the filter size and level l . The biases b_l are vectors of length n_l . The non-linearity function \mathcal{O} applies element-wise and is fixed. The last layer f^L has to convert the low resolution feature maps to a high resolution image I^{SR} .

In an embodiment, the high-resolution data is super resolved from the low-resolution feature maps by a sub-pixel convolution layer, to learn the upscaling operation for image and video super resolution. If the last or final layer performs upscaling, feature extraction occurs through non-linear convolutions in the low-resolution dimensions, so a smaller filter size can be used to integrate the same information relative to doing so in the high-resolution dimensions thus reducing computational complexity.

In an embodiment, the convolutional neural network consists of a plurality of non-linear mapping layers followed by a sub-pixel convolution layer. The benefit is the reduced computational complexity compared to a three-layer convolutional network used to represent the patch extraction and representation, non-linear mapping and reconstruction stages in the conventional sparse-coding-based image reconstruction methods and using an upsampled and interpolated version of the low-resolution input. The convolutional neu-

ral network of the embodiment uses learned filters on the feature maps to super resolved the low-resolution data into high-resolution data (instead of using hand crafted interpolation or a single filter on the input images/frames).

Instead of using a deconvolution layer to recover resolution from max-pooling and other image down-sampling layers, according to some embodiments upscaling of a low-resolution image is performed by convolution.

In an embodiment, convolution is performed with a fractional stride of

$$\frac{1}{r}$$

in the low-resolution space, which can be implemented by interpolation, perforate or un-pooling from low-resolution space to high-resolution space followed by a convolution with a stride of 1 in high-resolution space.

In another embodiment, convolution with a stride of

$$\frac{1}{r}$$

in the low-resolution space is performed with a filter W_d of size k_s with a weight spacing

$$\frac{1}{r}$$

to activate different parts of W_d for the convolution. The weights that fall between the pixels are not activated and therefore not calculated. The number of activation patterns is r^2 and each activation pattern, according to its location, has at most

$$\left(\frac{k_s}{2}\right)^2$$

weights activated. These patterns are periodically activated during the convolution of the filter across the image depending on different sub-pixel: $\text{mod}(x, r)$, $\text{mod}(y, r)$ where x, y are the output pixel coordinates in high-resolution space. In the situation where $\text{mod}(k_s, r k_s) = 0$, for a neural network with 1 layers, the low-resolution feature maps $f^{l-1}(I^{LR})$ from the last non-linear mapping layer are aggregated to produce the final estimate of the high-resolution image $f^L(I^{LR})$ using the following formula:

$$I^{SR} = f^L(I^{LR}) = \text{PS}(W_L * f^{L-1}(I^{LR}) + b_L)$$

where W_L and b_L represent learnable filters and biases respectively. The final convolution filter W_L is of size $C_{L-1} \times r^2 C \times k_L \times k_L$, where C is the number of colour channels in the original image and r the resampling ratio. PS is a periodic shuffling operator. Thus, after the convolution, but before applying PS, we have a tensor that has dimension $rH \times rW \times C \times r^2$, where h and v are horizontal and vertical dimensions of the high-resolution image. The result of applying PS is a $rH \times rW \times C$ array, i.e. the same dimensions

as the high-resolution image. This is achieved by periodic shuffling PS described in the following way:

$$PS(T)_{x,y,c} = T_{\left[\frac{x}{r}\right], \left[\frac{y}{r}\right], c \cdot r \cdot \text{mod}(y,r) + c \cdot \text{mod}(x,r)}$$

It is easy to see that when

$$k_L = \frac{k_s}{r}$$

and $\text{mod}(k_s, r) = 0$ it is equivalent to sub-pixel convolution in the low-resolution space with the filter W_d . This last layer, or sub-pixel convolution layer, produces a high resolution image from the low resolution feature maps directly, with a more distinguishable filter for each feature map and can operate with increased efficiency as the operation of the periodic shuffling described can be processed in parallel in a single cycle.

Given a training set consisting of high resolution image examples I_n^{HR} , $n=1 \dots N$, the corresponding low resolution images I_n^{LR} , $n=1 \dots N$ are generated, and the pixel-wise mean squared error (MSE) of the reconstruction is calculated as an objective function to train the network:

$$l(W_{1:L}, b_{1:L}) = \frac{1}{r^2 HW} \sum_{x=1}^{rH} \sum_{y=1}^{rW} (I_{x,y}^{HR} - f_{x,y}^L(I^{LR}))^2$$

Referring to FIG. 8, there is shown a network architecture **800** for super resolution using three input frames **810a**, **810b** and **810c** (and then **810b**, **810c** and **810d**). Only the input layer of the neural network is modified to jointly handle three input frames for the super resolution prediction of only the middle frame **830b** (or **830c**). The hidden layers **820a**, **820b** are identical to those shown in FIG. 7.

As shown in FIG. 8, in an embodiment, spatio-temporal convolutions are performed in the initial layer of the neural network. Using a spatio-temporal input layer as an extension of the network to handle video data allows reconstruction performance to be improved for video data by exploiting the temporal redundancy between consecutive frames in a video.

In some embodiments, the settings $l=3$, $(f_1, n_1)=(5, 64)$, $(f_2, n_2)=(3, 32)$, and $f_3=3$ can be used. In the training phase of these embodiments, $17r \times 17r$ pixel sub-images were randomly selected from the training ground truth images I^{HR} , where r is the upscaling factor. In other embodiments, other parameters can be chosen based on the available computing hardware and/or training time. To synthesise the low resolution samples I^{LR} , I^{HR} were blurred using a Gaussian filter in these embodiments and were sub-sampled by the upscaling factor. In these embodiments, the sub-images were extracted from original images with a stride of $(17 - \sum \text{mod}(f, 2)) \times r$ from I^{HR} and a stride of $(17 - \sum \text{mod}(f, 2))$ from I^{LR} . In other embodiments, arbitrary sized sub-images can be selected: smaller sub-images can enable faster updates in training (as updates can be made faster since fewer convolutions are required per sub-image) and can also allow more individual training examples to be used and can further allow an increase in representations; while larger sub-images can be more process-efficient for batch systems, particularly where there is sufficient training data available. It

can be noted that in most embodiments the lower limit on sub-image size is dependent on the filter sizes used in a neural network while the upper limit is memory- and data-set-related. This ensures that all pixels in the original image appear once and only once as the ground truth of the training data. tan h may be used as the activation function in these embodiments. In some embodiments, the training may stop after no improvement of the cost function is observed after 100 epochs. In some of these embodiments, the initial learning rate may be set to 0.01 and final learning rate may be set to 0.0001 and updated gradually when the improvement of the cost function is smaller than a threshold μ .

In some embodiments, the network is adapted for video super-resolution with a modification to the input layer. Instead of handling single frames independently, consecutive frames are jointly processed by feeding the consecutive multi-frame data such that the input can be expressed or converted into a three-dimensional vector of corresponding patches from multiple frames such that the temporal changes between consecutive frames are captured to the network. In an alternative embodiment, fixed temporal sampling can be used where videos have varying or different frame rates, optionally selecting the lowest frame rate when the network is trained. After patch extraction, the data is processed by hidden layers through the network and results in an output layer producing a single frame of high-resolution image corresponding to the network prediction of the middle frame in the input consecutive frames. In some alternative embodiments, there are both multiple input frame and multiple output frames.

In use for video processing, blocks of 3, 5 and 7 (or more) consecutive video frames may be used to super resolve the middle frame or frames. It should be noted that an increased number of frames could become impractical as each additional frame provides decreasingly less of a gain in accuracy. In order to handle the first and last frames of videos that do not have enough neighbouring frames to form a block of consecutive frames, the first and last frame may be repeated 1, 2 and 3 times respectively. The peak signal to noise ratio (PSNR) may be used as the performance metric to evaluate the models used.

Referring now to FIGS. 4 and 6, an embodiment for reconstructing the video encoded using the technique will now be described in detail.

First, the data **130** is received from the network. The data received **130** depends on how the data was prepared for transmission in step **120** (or **210**) as detailed above and will include the video data and one or more reconstruction models for that video data. It is possible that the video data and one or more reconstruction models are not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required to decode or reconstruct the higher-resolution video from the data transmitted over the network.

At steps **140** and **150** (and step **220**), the received data is prepared for reconstruction. This step generally involves separating the low-resolution video from the reconstruction models (step **220**). Optionally, the low-resolution video is decompressed using the video codec used for transmission into full-resolution image frames (step **230**) and each of the frames is matched up with the corresponding reconstruction model, which are also unpacked into frame-by-frame order (step **240**).

At step **160** (step **250**), the reconstruction model is applied to each of the frames to output higher-resolution frames. The reconstruction, or decoding, process involves applying the optimised super resolution convolutional neural network

model, or reconstruction model, for each scene in order to restore the lower-resolution video to its original resolution having substantially the same quality as the original high-resolution video. Given the corresponding models for each lower-resolution frame, the original higher-resolution frames can be reconstructed with high accuracy.

In a further embodiment, at least a portion of the process described within FIG. 3 can take place within a 'cloud computing' system. In this embodiment, original video data 70 is transmitted to an off-site computing system wherein one or more of the processes outlined in this specification take place. A different section of video data may be processed in parallel on the off-site computing system. The limit of the number of sections of video processed in parallel is dependent on the computational complexity of the sections of video and the processing power of the cloud servers. Once the cloud-based processing is complete, data can be transmitted back to a local computing system. The cloud service can be used as a real-time relay server, receiving lower-resolution videos, the resolutions of which are then to be increased before being re-transmitted. This arrangement can relay data from bandwidth-restrained locations to locations where there is not such a constraint on bandwidth.

In a further embodiment, the training process for the example based models takes place exclusively for the reconstruction of the high-frequency components of the higher-resolution section of video. The results may then be added as a residue to a section of video reconstructed using bi-cubic interpolation.

In a further embodiment, the architecture of any convolutional neural networks used in the reconstruction process is amended such that the upscaling, in terms of physically changing the number of pixels, occurs in the middle or the last layer of the network by way of a deconvolution function. The first layer of the network is described as obtaining low resolution features whilst the high resolution features are only learnt in the last layer. It is not necessary to learn the low resolution features in a high resolution space. By keeping the resolution low for the first couple of layers of convolutions, the number of operations required can be reduced. The computational performance then depends more on the resolution of the input rather than resolution of the output. Alternatively, the network layers can be reorganised such that the output of the sub-pixel convolution layer of the convolutional neural network (which can be the last layer or a preceding layer, depending on the embodiment) predicts the square of the upscaling factor number of pixels (e.g. r^2) instead of the number of upscaling factor pixels (e.g. r). The other pixels represent the neighbouring pixel to the original input pixel in low resolution. Every 4 pixels must then be reshaped and reordered to form the high resolution image. A higher video processing speed can be obtained using this approach whilst maintaining a high level of quality for the output. The number of convolutions required is therefore reduced. This embodiment may be used in conjunction with any other embodiment.

In some and other embodiments, a neural network can have multi-stage upscaling (or other function) where an earlier layer upscales and then a later layer upscales, for example a middle layer upscales by $2\times$ and then the last layer upscales by $2\times$. This type of "chained" approach can allow for neural networks to be trained in a long network (or "chain") of functional layers, for example having a range of upscaling factors (e.g. $2\times$, $3\times$ and $4\times$) with multiple upscaling layers and output layers. One of more of these layers can be sub-pixel convolution layers of the described embodiments.

The complexity of the reconstruction process is far simpler compared to that of training the model. The process is very similar to a single iteration of the training process however there is no optimisation required and it is simply applying the learnt convolutional neural network model to reconstruct the image on a patch-by-patch basis. The size of a single patch may vary, according to the resolution requirements of the output section of video. The patch size will be determined offline via hyper-parameter optimisations when training models. Generally, a larger patch size will reduce the computational requirements for processing a section of video, and also reduce the bandwidth required to transmit the section of video. When bandwidth is particularly limited, the processing power of an end user terminal is low, or a user is willing to endure a lower quality section of video, the patch size may be increased to allow for video transmission under one or more of those circumstances. The quality of the reconstruction will be compromised but the processing power required at the end-user terminal will be significantly reduced.

A singular model is applied to each scene and can be sent as accompanying data in the video container ahead of the scene in which it is used. The model associated with each scene can either be indexed in a data stream, synchronised with the video and audio stream, or indexed in a frame-to-model look-up table which can be sent ahead of the video stream.

At step 170, each of the segments of video are output from the reconstruction process as higher-resolution frames at the same resolution as the original video 70. The quality of the output video 180 is substantially similar to that of the original video 70, within the error bounds optionally applied by the machine learning process that develops the reconstruction model at step 110 (or step 200) of the encoding process.

At step 180, the segments of video are combined such that the video can be displayed. Optionally at step 170 or 180, the video can be re-encoded with the original codec used on the original video 70 or a more predetermined optimal codec, to enable playback on a decoder or display device.

The use of the final layer to perform super resolution can be used to perform upscaling for low-resolution images and/or video without knowledge of any original high-resolution content or where there is no higher-resolution content for the low-resolution images and/or video. Additionally, the use of the spatio-temporal approach can be used to perform upscaling for low-resolution images and/or video without knowledge of any original high-resolution content or where there is no higher-resolution content for the low-resolution images and/or video. Further, both the use of the final layer to perform super resolution and the use of the spatio-temporal approach can be used in combination to perform upscaling for low-resolution images and/or video without knowledge of any original high-resolution content or where there is no higher-resolution content for the low-resolution images and/or video. Specifically, some embodiments can use one or more generic example based models, for example selecting these from a library of generic example based models and for example applying some criteria in selecting such example based models, that use either or both of example based models having a final layer that performs super resolution and/or use a spatio-temporal approach in the input layer of the example based model.

Real Time Model Development/Selection

Referring now to FIG. 27, an embodiment using the technique to encode visual data will now be described in detail. These embodiments can be used in combination with

other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

Original video data **2740** can be provided into the method or system of the embodiment using a camera **2735** and usually comprises a high-resolution video, for example having a resolution of 1920 pixels by 1080 pixels (also known as “1080p” video) or 3840 pixels by 2160 pixels (also known as “4K” video). This video data can be encoded in a variety of known video codecs, such as H.264 or VP8, but can be any video data for which the system or method is able to decode into the component frames of the video depending on the embodiment.

In some embodiments, the original video data **2740** is then split into single full-resolution frames within step **150** i.e. into a sequence of images at the full resolution of the original video data **140**. In some of these embodiments, for some video codecs, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

In some embodiments, model selection or at least initial model selection can be performed based on the uncompressed video data.

Optionally, within step **2750**, in some embodiments the full-resolution frames can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection”. The video data is split or clustered into scenes to enable more specific optimisation. By scene, it is meant a consecutive group or sequence of frames comprising a section of video, which at the coarsest level can be the entire video or at the most granular level can be a single frame. In some embodiments, the scenes can be arranged according to the order in which they were filmed or into a sequence for appropriate decoding of any group of pictures (as scenes will typically comprise related visual content and appropriate ordering can allow compression to work efficiently, for example).

In some embodiments, regardless whether the video has been broken into frames or scenes (i.e. groups of multiple frames) in step **2750** or remains as a sequence of frames from step **2740**, each frame can be down-sampled into lower-resolution frames at a suitably low resolution. Optionally, in some embodiments this step can occur before the frames are grouped into scenes in step **2750**. In some embodiments, the lower-resolution frame is optionally 33% to 50% of the data size relative to the data size of the original-resolution frame, but can be any resolution that is lower than the original resolution of the video. In other embodiments, the quality can be reduced by quantisation and/or compression instead of reducing the resolution of the visual data or in addition to reducing the resolution of the visual data.

At step **2755**, a scene is selected in some embodiments. In this embodiment used for real-time encoding, this scene can initially be the first frame chronologically that was recorded by the camera **2735**, or the first frame appearing in the original video file **2740**. Following this frame, a second frame to be recorded or a second frame appearing in the original video file **2740** can next be selected. After that would appear a third frame, and so on, until the broadcast was complete.

A first example based model is taken from a library of example based models and analysed for use in the reconstruction of data in step **2760**. A second example based model can be taken and analysed in parallel in step **2770**, as can a third example based model as shown in step **2780**. There is no fixed upper limit to the variable n , variable n

being the number of example based models which can be analysed in parallel, where n is at least two.

In step **2795**, the most accurate example based model is chosen to be used for the reconstruction of data. The most accurate example based model can be defined as an example based model which would result in the highest quality of reconstruction, wherein quality can be defined using any of: an error rate; a peak signal-to-noise ratio; or a structural similarity index when compared to the original video file **2740**.

In step **210**, the lower-resolution scene and the example based model for the reconstruction of that scene are output for transmission over a network. Optionally, a reference to the model can be transmitted in lieu of the actual model. This requires a synchronised or matched library at both the transmitting and receiving nodes of the process, such that the reference can be used by the library at the transmitting node, and the model identified from the reference to the same model in the library at the receiving node as in the library at the transmitting node.

Optionally, the video and model (or model reference) can be stored together within a suitable data container format such as a Matroska Multimedia Container (otherwise known as a MKV container). Alternatively, the video and model can be combined with other sections, or the entire video and placed into a suitable data container format. At step **210** the low-resolution video frames can be re-encoded using either the original video codec applied to the original video data **2740** or, alternatively, a more optimal video codec can be applied to the video data to produce a smaller output file. Optionally, if scene detection or time stamping was performed, the data output for transmission can include either a list of scenes or time stamp data respectively, or this data could be stored within the data container.

There can be a number of variations to the encoding framework described above. For the above described technique, it can be assumed that there would be a one-to-one mapping between scene and example based model, however this does not need to hold true as each scene could be accompanied by more than one example based model.

Referring now to FIG. **6**, an embodiment for reconstructing the video encoded using the technique will now be described in detail.

First, the transmitted data **210** is received from the network. The data received **210** depends on how the data was prepared for transmission in step **2750** as detailed above and will include the video data and one or more reconstruction models for that video data. It is possible that the video data and one or more reconstruction models are not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required to decode or reconstruct the higher-resolution video from the data transmitted over the network.

At step **220**, the transmitted data is prepared for reconstruction. This step comprises separating the low-resolution video from the one or more example based models. Optionally, the low-resolution video is decompressed using the video codec used for transmission into full-resolution image frames and each of the frames is matched up with the corresponding reconstruction model, which are also unpacked into frame-by-frame order.

At step **250**, the reconstruction model is applied to each of the scenes to output higher-resolution scenes. The reconstruction, or decoding, process involves applying the optimised super resolution convolutional neural network model, or reconstruction model, for each scene in order to restore the lower-resolution video to its original resolution having

substantially the same quality as the original high-resolution video. Given the corresponding models for each lower-resolution scene, the original higher-resolution scenes can be reconstructed with high accuracy, where accuracy is defined as above.

This embodiment requires no optimisation, and comprises applying a previously learnt example based model to reconstruct a scene, and does not necessarily include any adaptation or modification of the existing example based models in the library. The one or more example based models associated with each scene can be indexed in a data stream, synchronised with the video and audio stream, or indexed in a frame-to-model look-up table that can be sent ahead of the video stream.

The image reconstruction has a linear computational complexity dependent on the total number of pixels of each frame, as it requires a series of convolution functions (typically less than 3) to be applied for each relevant image patch centred around each pixel. The basic reconstruction step for each image can be used with or without modification in some embodiments from the approach described in the paper “Learning a Deep Convolutional Network for Image Super-Resolution” by Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang published in D. Fleet et al. (Eds.): ECCV 2014, Part IV, LNCS 8692, pp. 184-199, 2014, however these functions can be applied in parallel across multiple images and it is feasible to perform real-time reconstruction for video playback.

Within step **250**, each of the segments of video is output from the reconstruction process as higher-resolution frames and at substantially the same resolution as the original video **2440**.

Further within step **250**, the segments of video are combined such that the video can be displayed. Optionally the video can be re-encoded with the original codec used on the original video **2740** or a predetermined more optimal codec, to enable playback on a decoder or display device.

In a further embodiment, at least a portion of the process described within FIG. **27** can take place within a ‘cloud computing’ system. In this embodiment, the original video file **2740** is transmitted to an off-site computing system, preferably a scalable computing system, where the first scene is prepared **2750**. As before, a plurality of example based models **2760**, **2770**, **2780**, **2790** can then be analysed for their use in reconstruction of the data, before a most accurate model **2795** is chosen and a transmitted package **210** prepared. The transmitted package **210** is then sent from the cloud server to a local server, where the reconstruction process **220**, **230**, **240**, **250** can take place. While the first scene is being processed in this way, a second section of video can be processed in parallel on the off-site computing system. The limit of the number of sections of video processed in parallel is dependent on the computational complexity of the sections of video and the processing power of the cloud servers. The cloud service can be used as a real-time relay server, receiving lower-resolution videos, the resolutions of which are then to be increased before being re-transmitted. This arrangement can relay data from bandwidth-restrained locations to locations where there is not such a constraint on bandwidth.

In a further embodiment, one or more sections of video are transmitted via the process outlined above, in advance of a live recording. Example based models that are likely to be useful can therefore be identified before they are required, and the library of models from which they are selected reduced to those more likely to be used. The selection procedure **2795** can therefore be made faster or with reduced

computation, as there are fewer models to choose between, hence the lag between the recording of an event and the higher-resolution reconstruction being displayed can be reduced compared to comparing the performance of all of the models in the library of models.

Model Libraries

Referring now to FIG. **25**, an embodiment using another technique to encode visual data will now be described in detail. These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

Original video data **2510** is provided into the method or system using the technique and is a high-resolution video, for example having a resolution of 1920 pixels by 1080 pixels (also known as “1080p” video) or 3840 pixels by 2160 pixels (also known as “4K” video). This video data can be encoded in a variety of known video codecs, such as H.264, VP8 or VP9 but can be any video data for which the system or method is able to decode into the component frames of the video.

The original video data **2510** is then split into single full-resolution frames at step **2520**, i.e. into a sequence of images at the full resolution of the original video data **2510**. For some video codecs, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

Optionally, at step **2520**, the full-resolution frames can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection”. The video data is split or clustered into scenes to enable more specific training and optimisation. By scene, it is meant a consecutive group or sequence of frames, which at the coarsest level can be the entire video or at the most granular level can be a single frame.

Regardless of whether the video has been broken into frames or scenes (i.e. groups of multiple frames) or remains as a sequence of frames, each frame is down-sampled into lower resolution frames at a suitably lower resolution for transmission over a network, for example the Internet. Optionally, this step can occur before the frames are grouped into scenes in step **2520**. The lower-resolution frame is optionally 25% to 50% or 10% to 50% of the data size relative to the data size of the original-resolution frame, but the lower resolution can be any resolution that is lower than the original resolution of the video. This will result in the full lower resolution video being smaller in size than the original-resolution video.

An analysis step **2530** is performed on the frames of the down-sampled video in order to find a reconstruction model from a library that can be used to increase the resolution of the scene and enhance the quality of the video. A model is selected based on one or more metrics of the selected scene that can be compared to metrics associated with the reconstruction models stored in the library, and an image enhancement procedure is performed on the down-sampled scene. The quality of the enhanced scene is compared with the original using objective metrics such as error rate, PSNR and SSIM and/or subjective measures. An appropriate model is then selected at step **240** based on these quality comparisons. The library from which the models are selected comprises a set of pre-trained models which have been generated from example, or training, videos and which are associated with metrics to enable comparison of the video from which the models were generated with the selected scene being enhanced. The library from which the model is selected may

be stored at one or more specific nodes in the network, or may be distributed over two or more nodes.

By employing a deep learning approach when creating the models, a non-linear hierarchical model can be selected to reconstruct a higher-resolution frame from the lower-resolution frame.

The non-linear models are more accurate in reconstructing higher-resolution frames than dictionary-based approaches. The non-linear models selected are small convolutional neural networks rather than over-complete dictionaries. In contrast to the local patch averaging approach that tends to be used in reconstruction by dictionary learning approaches, the use of a convolutional neural network model also allows a more appropriate filter to be selected for final reconstruction where neighbouring patches are considered, which can avoid unintended smoothing of the reconstructed higher-resolution image.

In step **2550**, the portion of the low-resolution video and a library reference for reconstruction model for that portion of video are output for transmission. Optionally, the video and library reference can be stored together within a suitable data container format such as a Matroska Multimedia Container (otherwise known as a MKV container). Alternatively, the video and library reference can be combined with other sections or the entire video and placed into a suitable data container format. At step **2550**, the low-resolution video frames can be re-encoded using either the original video codec applied to the original video data **2510** or, alternatively, a more optimal video codec can be applied to the video data to produce output video data. Optionally, if scene detection or time stamping was performed, the data output for transmission can include either a list of scenes or time stamp data respectively, or this data could be stored within the data container. Optionally, a list (or data stream) of model references that is synchronised with the output low resolution video frames can also be included. Alternatively, the transmission package can include a list of references to library content/models/parameters for each frame, which can be ordered by either frame playback or decode order, and this can be synchronised to the video frames. A further alternative is that the transmission package can include a list of scene transitions (at defined frame intervals) and then a list of references to library content/models/parameters for a matching number of scenes or transitions.

Referring now to FIG. **26**, an embodiment for reconstructing the video encoded using the technique will now be described in detail.

First, the data **2610** is received from the network. The data received **2610** depends on how the data was prepared for transmission in step **2550** as detailed above and will include the video data and one or more library references to reconstruction models for that video data. It is possible that the video data and one or more library references are not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required to decode or reconstruct the higher-resolution video from the data transmitted over the network.

At steps **2620** and **2630**, the received data is prepared for reconstruction. These steps generally involve separating the low-resolution video from the library references to reconstruction models. Optionally, the low-resolution video is decompressed using the video codec used for transmission into full-resolution image frames and each of the frames is matched up with the corresponding library reference to a reconstruction model, which are also unpacked into frame-by-frame order.

At step **2630**, the reconstruction model corresponding to the library reference is obtained from a stored library of reconstruction models and at step **2640** is applied to each of the frames to output higher-resolution frames. The reconstruction, or decoding, process involves applying the reconstruction model for each scene in order to restore the lower-resolution video to its original resolution having substantially the same quality as the original high-resolution video. Given the corresponding models for each lower-resolution frame, the original higher-resolution frames can be reconstructed with high accuracy.

The complexity of the reconstruction process is far simpler compared to that of training the model. The process, known as the feed-forward process or pass for neural networks, is very similar to a single iteration of the training process however there is no optimisation required and it is simply applying the selected convolutional neural network model to reconstruct the image on a patch-by-patch basis. In comparison, training neural networks can also incorporate a back-propagation process or pass, depending on embodiment. A singular model is applied to each scene and the library reference to it can be sent as accompanying data in the video container ahead of the scene in which it is used. The library reference associated with each scene can either be indexed in a data stream, synchronised with the video and audio stream, or indexed in a frame-to-model look-up table which can be sent ahead of the video stream.

At step **2650**, each of the segments of video are output from the reconstruction process as higher-resolution frames at the same resolution as the original video **210**. The quality of the output video is substantially similar to that of the original video **210**, within the error bounds optionally applied by the machine learning process that selects the reconstruction model at step **2650** of the encoding process.

At step **2650**, the segments of video are combined such that the video can be displayed. Optionally at step **2650**, the video can be re-encoded with the original codec used on the original video **2510** or a more predetermined optimal codec, to enable playback on a decoder or display device.

Upscaling & Visual Data Enhancement

Super resolution approaches are usually split into learning- or example-based approaches and interpolation-based (multi-frame) approaches. Some embodiments are concerned only with learning- or example-based approaches to super resolution. Specifically, these embodiments can work with most or all learning- or example-based techniques where there can be a set of different upscaling results depending on the selected parameters for the techniques.

In some embodiments, super resolution techniques can output representations that can be used to enhance the higher-resolution images created from lower-resolution images. To improve the effectiveness of these representations in some embodiments, learning- or example-based approaches incorporate machine learning. When using dictionary representations for images, this combination is generally referred to as dictionary learning. In dictionary learning, where sufficient representations are not available in an existing library of representations (or there is no library available), machine learning techniques are employed to tailor dictionary atoms such that they can adapt to the image features and obtain more accurate representations.

It is noted that atoms are not selected locally within the dictionary, but instead are chosen as the linear combination that best approximates the signal patch for a maximum number of atoms allowed and irrespective of their location within the dictionary. Without a constraint that the atoms

must be orthogonal to one another, larger dictionaries than the signal space that the dictionary is intended to represent can be created.

Embodiments can use dictionary learning reconstruction models or convolutional neural network reconstruction models for up-scaling, or a mixture of these two techniques. In some embodiments, a library of reconstruction models is stored that can be generated from example, or training, video data where both the original and reduced-resolution video can be compared. Along with the reconstruction models, in some embodiments data needs to be stored relating to the example or training video for each reconstruction model in the library to enable each model to be matched to a scene that is being up-scaled. In these embodiments the data stored relating to the example or training video can be metadata or metrics related to the video data, or it can be samples or features of the example or training video.

Referring now to FIG. 15, embodiments using a technique 1500 to increase the resolution of visual data will now be described in detail. These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

In these embodiments, received video data 1540 is provided into a decoder system and is a lower-resolution video encoded in a standard video format. This video format can be a variety of known video codecs, depending on the embodiment, and in embodiments can be codecs such as H.264 or VP8 or can be any visual data which the system is able to decode into component sections.

In embodiments, the system then separates video data 1510 into single frames at step 1520, i.e. into a sequence of images at the full resolution of the received video data 1510. For some video codecs, depending on the embodiment, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

Optionally, as shown in FIG. 16, in some embodiments the frames extracted in step 1520 from the video 1510 can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection” in step 1610. In these embodiments step 1610 involves identifying sequential scenes once these are detected. In some of these embodiments, the detected scenes can enable a more accurate reconstruction model or models to be selected from the library of reconstruction models as appropriate from the scene or frames within the scene in step 1620, which is a modified version of step 1530 in process 1500 where the most appropriate model is selected on a frame-by-frame basis. By scene, it is meant a consecutive group or sequence of frames, depending on embodiment, which at the coarsest level can be the entire video or at the most granular level can be a single section of a frame.

Alternatively, and in order for any delay to be reduced between the arrival of the lower resolution video 1510 and the generation of a higher resolution output video 1570, in some embodiments basic scene selection in step 1610 can be accomplished by grouping frames into scenes chronologically, for example by applying time stamps and collecting together a predetermined range of timestamps. Initially, in these embodiments the first frame or section of frame can be analysed and a metric created to enable reference to the library of reconstruction models in step 1620. In such embodiments, if a subsequent frame or section of a subsequent frame is sufficiently similar according to a comparison metric, then the subsequent frame or frames can be included

as part of a group of frames with the first frame. In these embodiments, this process can continue until a subsequent frame is not sufficiently similar to the previous frame or frames, at which point the frame or group of frames are collated into a scene in step 1610. In these embodiment the process then starts to find the next scene, starting from the insufficiently similar scene. In such embodiments, each scene is processed in the order in which it was decoded from the received video 1510.

Returning to the method 1500 shown in FIG. 15, in some embodiments an analysis step 1530 is performed on the frames of the received video in order to find a reconstruction model from a library in step 1540 which can be used to increase the resolution of the scene and enhance the quality of the video in step 1550, based on a metric of the selected scene that can be compared to metrics associated with the reconstruction models stored in the library. A model in step 1540 is selected, and an image enhancement procedure is performed in step 1550 using both the model selected in step 1540 on the selected frame generated in step 1520 from the received video 1510. The library from which the model is selected in step 1540 comprises a set of pre-trained models which have been generated from example, or training, videos and which are associated with metrics to enable comparison of the video from which the models were generated with a selected frame being enhanced in step 1550.

There are a number of methods for feature extraction approaches that can be used to create metrics for each frame and therefore which can be stored in the library in association with the pre-trained reconstruction models. In many computer vision applications, a common approach is to use a visual bag-of-words approach. Other popular methods include histogram extraction and scale invariant feature transform (SIFT) or GIST features. After feature extraction in step 1530, each scene from a target video can be matched to the “most similar” scene from a library in a number of ways in step 1540. The simplest approach is to use a pre-defined similarity or distance metric such as, Manhattan distance, Euclidean distance or structural similarity (SSIM) or additionally, a learned distance metric. A k-nearest neighbour (kNN) data-structure such as a ball-tree or a locality-sensitive hash table can be used to facilitate a direct nearest-neighbour search. Alternatively, each unique scene can be thought of as being “labelled”, then a variety of machine learning scene recognition and matching approaches can be used to classify new scenes into the “labelled” scenes of the library such as adversarial training techniques, where such techniques involve training another neural network to indicate whether the scene is labelled correctly.

In step 1560, the reconstructed scene for that section of video is optionally prepared for output 1510. Optionally at step 1560, the output video 1570 can be re-encoded with the original codec used on the original video 1510 or a more predetermined optimal codec, to enable playback on a decoder or display device.

It should be appreciated that the term ‘frame’, particularly in reference to grouping multiple frames into scenes, can refer to both an entire frame of a video and an area comprising a smaller section of a frame.

The video 1510 in process 1500 or 1600 can originate from physical media such as a DVD or other physical data storage medium, but can also originate from a local or remote streaming network service.

Referring now to FIG. 17, a modified process 1700 (based on process 1600 but which could equally be a modified version of process 1500) is shown.

Here, process **1600** in some embodiments is modified with further steps. Specifically, step **1710** begins with a high resolution video file, for example at HD resolution. Then, at step **1720**, the video is down-sampled to, for example, SD resolution to enable transmission over a bandwidth-restricted network. Optionally, in other embodiments, quantisation can be used to lower the bitrate of visual data such as video. At step **1730**, the lower-resolution or quantised video (depending on embodiment) is transmitted over a network becoming the input video data **1510** of either process **1500** or **1600** described above in relation to FIG. **15** or **16** respectively.

Further, output video **1570** in process **1500** or **1600** can be output directly to a display or may be stored for viewing on a display on a local or remote storage device, or forwarded to a remote node for storage or viewing as required.

In alternative embodiments, the input video concerned may be media for playback, such as recorded video or live streamed video, or it can be videoconference video or any other video source such as video recorded or being recorded on a portable device such as a mobile phone or a video recording device such as a video camera or surveillance camera.

It should also be appreciated that particular combinations of the various features described and defined in any aspects of disclosed implementations can be implemented and/or supplied and/or used independently.

Specific Description—Machine Learning Optimisation

Embodiments relating to a dictionary learning approach with dictionary-based models and embodiments relating to a deep learning approach using convolutional neural network models will now be described. Other suitable machine learning techniques will also be able to benefit from using a library of pre-trained model. These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

Referring now to FIG. **9**, an embodiment of a machine learning process will now be described in detail.

Data **910** is provided to the method or system **900**, and may for example be in the form of a video data file or image data. Other data formats may also be used. Specific metrics are extracted from, or generated based on, the data in step **920**. The metrics can be based on measuring which of the pre-trained models generates the highest quality output data. Quality can be defined using any or a combination of: an error rate; a peak signal-to-noise ratio; or a structural similarity index. These metrics are then used to select a pre-trained model from a library **942** in step **930**. The selected pre-trained model is then developed in step **940** so as to be more accurately able to perform the desired output function on the input data **910**. The output model **950** can then be used on the input data **910** either immediately or in future.

To generate the initial pre-trained models, the same process **900** can be used but instead of inputting live data **910**, training or example data is used.

The developed model is then saved in the library **942** in step **941**, to be used if required in the future. A model does therefore not have to be redeveloped if similar data is input in the future. The library **942** therefore grows as the system is used.

If employing a deep learning approach to generating the model in step **940**, a non-linear hierarchical model **950** can be created.

The training and optimisation process **900** can be configured according to a desired trade-off between computational time spent and desired quality of results. In general, the

number of iterations used during the training process **940** yields approximately logarithmic gains in reconstruction accuracy, so it is preferred to use an automatic threshold to stop further optimisation. When favouring quality of results, the automatic threshold can be set to a predetermined value of reconstruction error, for example by calculating the mean squared error, but other methods can also be used. Alternatively, the automatic threshold can be set to limit the training and optimisation process to a predetermined number of iterations. As a further alternative, a combination of these two factors can be used.

This process **900** can be of use in the fields of image and video enhancement. If the input data **910** is in the form of a section (frame or frames) of video, the selection of an initial model from the library **930** can be based on metrics associated with the style of video being processed, for example a fast-moving action scene, or a stationary background scene. The model **950** can then be developed in step **940** to more closely represent the section of video for which it is to be used. The training and optimisation process **900** can also be considered in the case of image registration and alignment. An initialisation can be provided which is closer to the most optimal image alignment with regards to a metric, for example a distance metric cost function.

Specific Description—Enhanced Libraries

Referring now to FIG. **10**, an embodiment using the technique to encode visual data will now be described in detail. These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

Original video data **1010** is provided into the method or system using the technique and is a high-resolution video, for example having a resolution of 1920 pixels by 1080 pixels (also known as “1080p” video) or 3840 pixels by 2160 pixels (also known as “4K” video). This video data can be encoded in a variety of known video codecs, such as H.264 or VP9 but can be any video data for which the system or method is able to decode into the component frames of the video.

The original video data **1010** is then split into single full-resolution frames at step **1020**, i.e. into a sequence of images at the full resolution of the original video data **210**. For some video codecs, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

Optionally, at step **1020**, the full-resolution frames can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection”. The video data is split or clustered into scenes to enable more specific training and optimisation. By scene, it is meant a consecutive group or sequence of frames, which at the coarsest level can be the entire video or at the most granular level can be a single frame.

Regardless whether the video has been broken into frames or scenes (i.e. groups of multiple frames) in step **1020** or remains as a sequence of frames, each frame is either down-sampled into lower resolution frames at a suitably lower resolution, or reproduced as a set of representations of the frame such as a matrix of vectors with generalised information for portions of the frame or groups of one or more pixels. Optionally, this step can occur before the frames are grouped into scenes. The lower-resolution frame is preferably 33% to 50% of the data size relative to the data size of the original-resolution frame, while the representations of the frame can be anything from 1% to 50% of the data size of the original-resolution frame. The lower reso-

lution frame can have any resolution that is lower than the original resolution of the video.

In step **1030**, an analysis is performed on frames of the downsampled video, or on the representations, or on the models alone in order to find a reconstruction model from a library of models which can be used to either recreate a close approximation of the high-resolution video purely from a model, or recreate a close approximation of the high-resolution video from a set of representations of the high-resolution video, or increase the resolution of the downsampled video and enhance the quality of the video. An initial model is selected based on a metric of the selected scene that can be compared to metrics associated with the reconstruction models stored in the library. The selection of the initial model may be based on only these metrics, or alternatively multiple initial models may be applied independently to the downsampled video to produce an enhanced lower-resolution frame or recreate a frame, the quality of the enhanced or recreated frame being compared with the original scene to select the most appropriate initial model from the group. The quality of the recreated or enhanced scene is compared with the original using objective metrics such as error rate, PSNR and SSIM and/or subjective measures. An appropriate model is then selected at step **1040** based on these quality comparisons as well as whether to use solely a model, or a set of representations of the frame, or the lower-resolution frame. The library from which the models are selected comprises a set of pre-trained models which have been generated from example, or training, videos and which are associated with metrics to enable comparison of the video from which the models were generated with the selected scene being enhanced.

In step **1060**, super resolution techniques are employed to create a specific model using the model selected at step **1040** as the initial model and using machine learning for each frame, so that the model can be used to substantially recreate the original resolution version of a lower-resolution frame, or a set of representations, or without use of either the original or lower-resolution frame or representations, and trained using machine learning based on knowledge of the original-resolution frame. This step is termed the training and optimisation process.

At step **1070**, the optimised reconstruction model is saved to the library for future use as an initial model in step **1030**.

Alternatively, the modified model can be developed from scratch, without using a known model as a starting point. A known model can then be compared with the generated model to produce a list of modifications required to produce the modified model from the known one.

In step **1050**, the portion of the low-resolution video, or the representations and/or the reconstruction model for that frame or portion of video are output for transmission. Rather than transmit the full reconstruction model, a library reference to the initial model selected at step **1040** can be output, together with the modifications required to produce the reconstruction model from it. Alternatively, instead of the lower-resolution video, transmitting a representation of the image can further reduce the data transmitted. As a further alternative, only transmitting a model still further reduces the data transmitted. Optionally, the video or representations and model can be stored together within a suitable data container format such as a Matroska Multimedia Container (otherwise known as a MKV container). Alternatively, the video or representations and model can be combined with other sections or the entire video and placed into a suitable data container format. At step **250**, the low-resolution video frames can be re-encoded using either the original video

codec applied to the original video data **1010** or, alternatively, a more optimal video codec can be applied to the video data to produce output video data **1050**. Optionally, if scene detection or time stamping was performed, the data output for transmission can include either a list of scenes or time stamp data respectively, or this data could be stored within the data container. Alternatively, a list of frame by frame references (which are ordered by either frame playback or decode order) can be used which is synced to the video frames. Another alternative is to include a list of scene transitions (as defined by frame numbers) along with an ordered list of references matching the number of scenes. The skilled person will appreciate that other examples are possible.

Referring now to FIG. **11**, an embodiment for reconstructing the video encoded using the technique will now be described in detail.

First, the data **1100** is received from the network. The data received **1100** depends on how the data was prepared for transmission in step **1050** as detailed above and will include the video data or representations of the video data and/or references to one or more initial reconstruction models for to recreate high-resolution video data, along with any modifications to the referenced models required to reproduce the final reconstruction models. It is possible that the video data or representations, references to the one or more initial reconstruction models and the modifications to those models are not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required to decode or reconstruct the higher-resolution video from the data transmitted over the network.

At steps **1120** and **1130**, the received data is prepared for reconstruction. This step generally involves separating the low-resolution video or representations if transmitted from the references to known reconstruction models and the modifications to those models required to reproduce the optimised reconstruction model. Optionally, the low-resolution video is decompressed at step **1120** using the video codec used for transmission into full-resolution image frames and each of the frames is matched up with the corresponding reconstruction model, which are also unpacked into frame-by-frame order.

The reconstruction models developed at step **1170** are recreated from the reference to a known reconstruction model and the modifications received at step **1110**. At step **1130** the received reference is used to locate the known reconstruction model from which the optimised reconstruction model was developed in a library of models stored at the second node. The received modifications are then applied to this known reconstruction model at step **1160**, recreating the optimised reconstruction model. After this reconstruction has taken place, the optimised reconstruction model is saved in the library at the second node for future reference at step **1170** using the same reference as was used when saving the model to the first node's library at step **1170**. An incremental (or decremental) counter can be used to create new references for the modified algorithm that would then be the same for both nodes (in a distributed system). Alternatively this could be performed via a centralised database.

At step **1140**, the modified reconstruction model is applied to each of the frames or representations if transmitted to output higher-resolution frames, else are used to reproduce the higher-resolution frames without a low resolution frame or representations. The reconstruction, or decoding, process involves applying the optimised super resolution convolutional neural network model, or reconstruction model, for each scene in order to recreate higher

resolution video having substantially the same quality as the original high-resolution video. Given the corresponding models and if transmitted each lower-resolution frame or set of representations, a higher resolution frame that is substantially visually identical to the original higher-resolution frames can be reconstructed with high accuracy.

At step **1150**, each of the segments of video are output from the reconstruction process as higher-resolution frames at the same resolution as the original video **210**. The quality of the output video **1150** is substantially similar to that of the original video **210**, within the error bounds optionally applied by the machine learning process that develops the reconstruction model at step **1060** of the encoding process.

At step **1140**, the segments of video are combined such that the video can be displayed. Optionally at step **1140**, the video can be re-encoded with the original codec used on the original video **1010** or a more predetermined optimal codec, to enable playback on a decoder or display device.

A modified embodiment of the above-described method is shown in FIGS. **4** and **5**. In this embodiment, the down-sampled video is not transmitted from the first node to the second node. Instead, any data required to produce the reconstructed video can be included in the transmitted modifications to a known reconstruction model. It is possible that no data is required and that only the use of a model is required to reconstruct substantially the high-resolution video. The second node reconstructs the video using only the modified reconstruction model obtained by applying the modifications to the known reconstruction model referenced in the transmission and, if required, the data required to produce the reconstructed video. This method can advantageously be used to reduce further the amount of data transmitted across the network, reducing the amount of data required to be transmitted.

FIG. **12** shows a further method for encoding video. The initial steps proceed as in FIG. **10**. The higher resolution video **1210** is prepared by separating it into scenes and down sampling at step **1220**. An analysis of each scene is performed in step **1230** to determine which known models will be used as a starting point for the model optimisation process. One or more appropriate models are selected at step **1240** based on a metric of each scene or frame that can be compared to metrics associated with the reconstruction models stored in the library. The selected known models are optimised at step **1260** using machine learning. However, instead of being trained to reproduce the higher resolution video from the downsampled one, they are instead trained to substantially reproduce the higher resolution video **1210** using only the optimised model itself or the optimised model and any data that can be incorporated into the model to seed recreation of the video frame or scene. A reference to the starting known reconstruction model together with the modifications required to reproduce the optimised reconstruction model for each scene or frame are transmitted to the second node at step **1260**, along with a list of scenes if known.

FIG. **13** shows a further method of decoding the received information at the second node to reproduce substantially the higher resolution video. The package transmitted at step **1260** of FIG. **12** is received by the second node at step **1310**, and unpacked at step **1320**. The references to known models are used to locate the known reconstruction models stored in a library at the second node at step **1330**. The modifications corresponding to these known reconstruction models are then applied at step **1340** to reproduce the optimised reconstruction model generated by the machine learning process in step **1260** of FIG. **12**. This model is optimised to substantially reproduce the original higher resolution video

1260 without the need for it to be applied to a corresponding low resolution video sample. Hence the second node can substantially reconstruct the original high resolution video from only the modified reconstruction model. This reconstruction is performed at step **1350**, and the resulting reconstructed higher resolution video is output at step **1360**.

Specific Description—Artefact Removal

To improve on the above mentioned approaches for image artefact removal, it is proposed to use deep learning techniques and neural networks such as recurrent neural network and convolutional neural network models.

In some cases fully connected neural networks can't scale up to larger sizes of network easily as the computational complexity soon becomes too great as the size of the network scales, but this depends on the application of the neural network and also other factors such as the kernel and filter sizes.

Neural network models can be transmitted along with the low-resolution and/or low-bitrate and/or low quality frames of video data because size of a convolutional neural network model or recurrent neural network model data is small enough to do so compared with the size of a learned dictionary for the same level of accuracy. In comparison, the data size of learned dictionaries means that it is impractical for these dictionaries to be transmitting along with the low-resolution images, especially in where the dictionaries are learned over-complete dictionaries

FIG. **22** illustrates an overview of embodiments of a method of generating models for use in image artefact removal. These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

In some embodiments, training image or video file is selected at step **2210** from a library of training images or videos stored at a network node and divided up into frames, sub-images or sequences of images (collectively referred to herein as scenes) at step **2220** depending on its content. In these embodiments, the scene type can then be classified into a particular category depending on its content using a metric.

In some embodiments, at step **2230**, image artefacts are purposefully generated in the scene by applying aggressive compression and/or quantisation algorithms to the scene, with the level of expected artefacts being controlled by the compression/quantisation level or quality level. This creates a training scene. In some of these embodiments an artefact grading technique is applied to this training scene at step **2240** to quantify the severity of the induced artefacts.

At step **2250**, in some embodiments, an image artefact removal model is trained on the training scene using machine learning techniques. In these embodiments the image artefact removal model can be generated such that it substantially cleans the training scene of image artefacts. In some of these embodiments, image quality metrics are used to compare the "cleaned" scene with the original, and the training process can be continued until the comparison produces results within some pre-defined threshold.

In some embodiments, the resulting image artefact removal model is then saved at step **2260**, along with metric data relating to the scene type classification and an artefact grade relating to the types and/or severity of the artefacts in the training scene.

The process can be repeated multiple times, depending on the embodiment, on the same training scene with different levels of compression and quantisation in order to train a set of models for different levels of artefact severity.

The process can also be repeated in some embodiments with different training scenes in order to generate a library of different image artefact removal models indexed by both their content type and artefact severity. In these embodiments a “matrix” of library models can be created, where one axis of the matrix can be thought of as the content type, and the other the artefact severity.

Each step of the above method, as used in some embodiments, will now be described in further detail.

The division of the training image or video into a scene or scenes at step 2220 can be achieved in several different ways in different embodiments. The training image or video may already be divided into scenes prior to being stored in the library of training images or videos in certain embodiments, hence this step of the method is not always necessary. For the training of the image artefact removal model, in some embodiments only one scene is used for each model generated. In these embodiments a single training image or video may therefore be used to train multiple different image artefact removal models relating to different content types present within the image or video.

Considering first the case of video data, in some embodiments the library video data can be split into single full-resolution frames, i.e. into a sequence of images at the full resolution of the original video data. For some video codecs, in certain embodiments, this will involve “uncompressing” or restoring the video data as, for example, common video compression techniques remove redundant (non-changing) features from sequential frames.

Optionally, in other embodiments the full-resolution frames can be grouped into scenes or sections of frames having common features, otherwise known as “scene selection”. In such embodiments, the video data can be split or clustered into scenes to enable more specific training and optimisation. By scene, it is meant a consecutive group or sequence of frames, which at the coarsest level can be the entire video or at the most granular level can be a single frame or portion of a frame depending on embodiment.

For the training of the image artefact removal model, in some embodiments only one scene is used for each model generated. A single training image or video may therefore be used in such embodiments to train multiple different image artefact removal models relating to different content types.

Once a scene has been generated, in some embodiments it is classified using metrics relating to properties of the scene. In embodiments, metrics that can be used for this classification can include probability or distance, depending on the classifier type or clustering used in each embodiment, and some function of the error or differences can be included in the metric depending on embodiment. Example metrics can include the differences between pixels or histogram differences such as the sum of absolute differences or Euclidean distance in these embodiments. Example metrics can also include mean squared error and peak signal-to-noise ratio in some embodiments. Alternatively, in other embodiments classification can be performed according to predetermined trained classifications, for example there could exist predetermined trained classifications for a variety of scenes or objects such as outdoors, indoors, mountains, bike, etc. The probability would be determined by the classifier used in the embodiment, for example embodiments could use a statistical similarity measure such as the Kullback-Leibler divergence. In other embodiments, a further alternative would be to use machine learning to develop a metric to be applied to the image-space, which would determine a formula for the distances. This classification using metrics is used to distinguish between scenes contain-

ing different content in at least some embodiments. The metric data determined in this way effectively assigns the scene a content type in such embodiments.

After the scene has been generated and classified, in at least some embodiments it is compressed and/or quantised at step 2230. In these embodiments, the compression of the scene can be implemented using any well-known lossy image/video compression algorithm. In such embodiments, the quantisation methods may include both colour quantisation, using for example a median cut algorithm or a clustering algorithm, and frequency quantisation.

In embodiments, the compression or quantisation process introduces image artefacts to the scene, producing a training scene that will be used together with the original uncompressed and unquantised version of the scene to train and/or optimise an image artefact removal model. In these embodiments, different levels of artefact severity can be introduced by varying the level of compression and/or quantisation performed at this stage.

Before the training and optimisation process begins in some embodiments, the artefact severity of the training scene is determined at step 2240 by using an artefact grading algorithm. Examples of such algorithms that can be used in embodiments are found in the paper “No-Reference Image Quality Assessment in the Spatial Domain” by Anish Mittal, Anush Krishna Moorthy and Alan Conrad Bovik published in IEEE TRANSACTIONS ON IMAGE PROCESSING, Vol. 21, No. 12, December 2012, which is incorporated herein by reference. By grading the artefact severity of a training scene in embodiments, image artefact removal models can be trained based on both scene content type and the severity of the artefacts present.

At step 250 the model training and optimisation process is performed on the training scene in some embodiments. Super resolution techniques and machine learning techniques are employed in some embodiments to create an image artefact removal model, such that the model can be used to substantially recreate the original “clean” version of the training scene. By employing a deep learning approach to generating the model in at least embodiments, a non-linear hierarchical model can be created to reconstruct the original clean scene from the compressed and/or quantised scene.

An example of a deep learning approach, but in respect of only still images and in relation to jpeg compression artefacts (but not, for example, other video artefacts such as motion blur and inter-frame artefacts), but which can be used with or without modification in some embodiments, is described in the paper “Compression Artifacts Reduction by a Deep Convolutional Network” by Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang published as arXiv:1504.06993v1 [cs.CV] 27 Apr. 2015 and this paper is incorporated herein by reference. This paper is related to generating an improved deep Super Resolution Convolutional Neural Network for image artefact removal from a shallow network model.

In some of the embodiments, the quality of the recreated or cleaned scene can be compared with the original using objective metrics such as error rate, PSNR and SSIM and/or subjective measures. In some of these embodiments, if the quality is found to be within some pre-defined threshold, then the model can be saved along with results from a classifier or outcomes identifying the scene content and the artefact severity that the model relates to. Otherwise, in some embodiments the training process continues until the predefined threshold is met. Alternatively or additionally, in some embodiments a pre-defined number of iterations of the

machine learning process can be used to limit the computational time spent training the model.

In some embodiments, the image artefact correction model trained in this way may be one of a non-linear hierarchical algorithm, a convolutional neural network, a recurrent neural network, a multi-layer (neural) network, an artificial neural network, a deep belief network, a dictionary learning algorithm, a parameter, a mapping function, a multi-layer feed-forward network, a non-linear filter, a dictionary, set or series of parameters or functions, or a series of functions.

In some embodiments, the original training scene generated in the scene selection process is used to generate multiple models for the same scene content type by repeating the compression and quantisation process with different levels of compression and/or quantisation, or by using different compression and/or quantisation algorithms. In these embodiments, this approach introduces a different level of artefact severity to the scene. Further, in some of these embodiments the training and optimisation process can then be repeated to generate a new image artefact removal model for the same content, but a different level of artefact severity.

In some embodiments, the trained and/or optimised image artefact removal models are saved to a library of image artefact removal models. In some of these embodiments, each model in the library is associated with a classifier identifying the scene content and the artefact severity that the model was trained on. In these embodiments, the library can therefore be thought of as a “matrix” of example based models, with its rows corresponding to the scene content type, and its columns corresponding to the image artefact severity.

The library can be stored on a node within a network in certain embodiments.

Referring now to FIG. 23, a method of using the image artefact removal models will now be described according to some embodiments.

A section of video or an image is received from a network at a receiving network node at step 2310 in some embodiments. In these embodiments, the data received depends on how the data was prepared for transmission at a transmitting network node and will include the video/image data or representations of the video/image data. It is possible in some of these embodiments that video data or representations is not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required before the artefact removal process can take place.

In some embodiments, at step 2320, the received video or image data is unpacked and divided into scenes depending on its content using the process described above in relation to the image artefact removal process. In some of these embodiments, the scenes are then classified at step 2330 using metrics to determine the scene content type.

Alternatively, in other embodiments the received video/image data can already have been divided into scenes and classified by the transmitting network node, in which case the data received by the receiving node will contain meta-data for each of the scenes contained within it that identifies the scene content type.

In either case, in some embodiments, at step 2330, the received scenes are subjected to an artefact grading algorithm or classifier to determine the level of artefact severity present in the scene. In some of these embodiments the source of these artefacts can be the video/image compression and quantisation process performed by the transmitting

node prior to transmission of the video/image data, or they could be introduced by faulty or lossy network transmission.

In some embodiments, the scene content type and image artefact severity are used to select the image artefact removal model from the library of image artefact removal models generated in the model creation process that best matches the scene content type and artefact severity of the received scene. In these embodiments, this matching may be performed at the receiving network node if the library is stored there. Alternatively, in other embodiments, if the library is stored remotely from the receiving network node, a request for the most suitable image artefact removal model may be transmitted from the receiving node to the location of the model library. In some of these embodiments the model library will then transmit the relevant image artefact removal model to the receiving network node. If no suitable image artefact removal is stored in the library then a generic model may be used instead. In such embodiments, a suitable model may be determined, for example, by requiring that the metric data relating to the content type and the artefact severity of the received video or image data lie within some predefined range of at least one of the image artefact removal models stored in the library.

In some embodiments, the image artefact removal model identified to be the most suitable (or the generic model in the case that no suitable model is present in the library or if the generic model is the most suitable model) is then applied to the received scene at step 2350 in order to substantially recreate the original video/image file largely free from image artefacts. In such embodiments, this fidelity correction process can result in a clean reconstructed video/image, which is then output by the receiving node at step 2360.

Referring now to FIG. 24, an alternative method of using the image artefact removal models according to some embodiments will now be described.

In such embodiments, this method relates to the combined use of image artefact removal and super resolution techniques to reconstruct a video or image from a downsampled (i.e. lower resolution than an original higher resolution) video or image data received from a network.

In some embodiments, firstly, at step 2410, the data package is received from the network. In these embodiments, the data received in the package depends on how the data was prepared for transmission by the transmitting network node and can include downsampled video or image data or representations of the video or image data and/or details of one or more reconstruction models for recreating high-resolution video data from the downsampled video or image data. In such embodiments, these details may be the reconstruction models themselves, or references to known reconstruction models stored at the receiving node or on the network.

In some embodiments, it is possible that the video data or representations, references to the one or more initial reconstruction models and the modifications to those models are not transmitted and/or received simultaneously so some buffering may be required to wait for all of the components required to decode or reconstruct the higher-resolution video from the data transmitted over the network.

Next, at step 2420, the received data is prepared for reconstruction in some embodiments. In these embodiments, this step generally involves separating the low-resolution video/image or representations if transmitted from the details of the relevant reconstruction models. Optionally, in other embodiments, the low-resolution video or image is decompressed using the video or image codec used for transmission into full-resolution scenes and each of the

frames is matched up with the corresponding reconstruction model, which are also unpacked into scene-by-scene order.

At step **2430**, in some embodiments, if the scene content has not already been determined by the transmitting node, then it is now determined by the receiving node.

In some embodiments, the unpacked video or image scenes are then graded to determine their artefact severity using an artefact grading algorithm.

In some embodiments, at step **2440**, the scene content type and image artefact severity are used to select the image artefact removal model from the library of image artefact removal models generated in the model creation process that best matches the scene content type and artefact severity of the received scene. In such embodiments, this matching can be performed at the receiving network node if the library is stored there. Alternatively, in other embodiments where the library is stored remotely from the receiving network node, a request for the most suitable image artefact removal model can be transmitted from the receiving node to the location of the model library. In such embodiments the model library will then transmit the relevant image artefact removal model to the receiving network node. If no suitable image artefact removal model is stored in the library then a generic model may be used instead in some of these embodiments.

In some embodiments, the image artefact removal model identified for each scene is then applied to the downsampled scene at step **2450** in order to substantially recreate the original downsampled image transmitted by the transmitting node. Alternatively, in other embodiments the identified model can be used after the upscaling process of step **2460** to substantially remove the upscaled image artefacts from the recreated higher resolution video.

At step **2460**, in some embodiments the relevant reconstruction models are then applied to each of the scenes to output higher-resolution scenes. In such embodiments the reconstruction, or decoding, process can involve applying the optimised super resolution convolutional neural network model, or reconstruction model, for each scene in order to recreate higher resolution video or image having substantially the same quality as the original high-resolution video or image from which the downsampled video or image was generated by the transmitting node. Given the corresponding reconstruction models and each lower-resolution scene or set of representations, in some embodiments a higher resolution frame that is substantially visually identical to the original higher-resolution frames can be reconstructed with high accuracy.

In some embodiments, the reconstructed higher resolution scene is then output by the receiving node at step **470**.

The above-described process is performed as a post-process to the standard decompression process used at the decoding node for received transmissions of, for example, video after the conventional decoding step and therefore “outside” the traditional compression pipeline in at least some embodiments.

Alternatively, in other embodiments, the image artefact removal process can be performed as part of the upscaling process itself. In such embodiments, several reconstruction models can be trained to reproduce the higher resolution image or video at the first node from a number of different downsampled images or videos, each with a different artefact severity. In such embodiments, these can either all be transmitted with the downsampled video, or the required model can be transmitted to the second network node from the first node once a request for the model containing the artefact severity of the received downsampled image or video has been sent from the second node to the first node.

In either case, in some embodiments the model best matching the artefact severity of the received downsampled image or video is used to substantially recreate the original high resolution video.

5 Spatio-Temporal Interpolation

In some embodiments, visual data being received may be of a lower than desirable quality, such as at a low resolution or a low frame rate. It will be appreciated other features may render the visual data lower than desirable quality. One or more hierarchical algorithms may be used to increase the quality of the visual data to a more desirable quality. For example, in some embodiments, the hierarchical algorithm may increase the resolution of the received visual data section or sections. In another embodiment the hierarchical algorithm may develop intermediate sections, such as a section to be placed between two received sections of visual data. In yet another embodiment, the hierarchical algorithm may be used to increase the quality of the received visual data by combining the above mentioned embodiments.

20 These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

In some embodiments, one or more hierarchical algorithms may be used to estimate higher-resolution versions of received lower-resolution frames. In such embodiments, upon receiving multiple lower-resolution frames, a hierarchical algorithm can be used to estimate a higher resolution version of a particular received frame. In some embodiments, on the receipt of three or more consecutive frames of lower-resolution visual data a hierarchical algorithm may be used to estimate a higher resolution version of one of the middle frames. In such embodiments, this estimation can be based not only on the lower-resolution version of the received frame, but also the lower-resolution version of the previous and subsequent frames also received. Alternatively, in other embodiments the hierarchical algorithm can be used to determine an unknown intermediate frame at the same resolution as previous and subsequent received frames.

Alternatively, in some other embodiments a hierarchical algorithm may be used to estimate higher-resolution versions of the received lower-resolution frames, as well as a higher-resolution version of an unknown intermediate frame. Furthermore, in some embodiments the hierarchical algorithm may also be used to estimate an unknown future or past frame based upon the received lower-resolution frames.

Online Training

In some embodiments, new visual data being transmitted from a node to one or more other nodes has no hierarchical algorithms trained to enhance a lower-quality version of the new visual data. For example, in some embodiments this new visual data may be live broadcast video data and in some embodiments the new visual data may be streamed or rendered e-gaming video content. Where there exist no hierarchical algorithms for this new visual data, specific or generic hierarchical algorithms need to be trained in order to enhance lower-quality versions of the new visual data in some of these embodiments and as per some of the embodiments described elsewhere.

60 In some embodiments, multiple hierarchical algorithms can be trained in parallel on the same visual data. In such embodiments, this allows for wider range of hierarchical algorithms to be explored within a given timeframe. In some of these embodiments, the most suitable algorithm from these developed algorithms can be selected for transmission across the network with the visual data. In some embodiments, developing the hierarchical algorithms in parallel can

be useful in situations where there is only a limited time available to develop the hierarchical algorithm.

In parallel to the development of hierarchical algorithms (whether singly or multiple in parallel), in some embodiments the visual data can be encoded in an encoder in preparation for its transmission across the network. In this way, the time taken to prepare both the hierarchical algorithm and the visual data for transmission can be reduced when compared to developing the hierarchical algorithm and encoding the visual data in series in such embodiments.

Once generic or specific hierarchical algorithms have been developed or trained, for example for a new game being streamed as e-gaming video content, these can be used to enhance the quality of lower-quality streamed visual data at receiving nodes.

These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification

Offline Training

In some embodiments, new visual data to be transmitted from a node to one or more other nodes has no hierarchical algorithms trained to enhance a lower-quality version of the new visual data. For example, in some embodiments this new visual data may be new video data such as a new film or television series and in some embodiments the new visual data may be a computer game that will in the future be used to generate streamed e-gaming video content. Where there exist no hierarchical algorithms for this new visual data, specific or generic hierarchical algorithms need to be trained in order to enhance lower-quality versions of the new visual data in some of these embodiments and as per some of the embodiments described elsewhere.

Often, the visual data will be encoded prior to transmission across the network using an encoder in certain embodiments. The encoding process may introduce encoder specific artefacts to the visual data in such embodiments. By refining the hierarchical algorithm by performing further training steps on the encoded video in some embodiments, the hierarchical algorithm can be trained to substantially correct image artefacts introduced by the encoder.

In some embodiments the hierarchical algorithms can be trained using training data corresponding to the exact visual data that will be transmitted across the network. In such embodiments, this approach can be particularly useful in situations where the visual data is known in advance, and where the visual data is likely to be transmitted across the network multiple times. For example, in some embodiments the hierarchical models can be trained on sections of an episode of a TV program that will be made available on an on-demand streaming service. In such embodiments the models trained on that particular episode can be transmitted alongside the lower-quality visual data and be used to enhance the lower-quality visual data to a higher-quality version of that episode.

Alternatively, in other embodiments the hierarchical algorithms can be trained on training data that is similar, but not identical, to the expected visual data on which they are to be used. In such embodiment these trained models can be associated with metric data relating to properties of the visual data on which they were trained. In some of these embodiments, this metric data can later be used to select an appropriate hierarchical algorithm to use to enhance visual data for which no specific model has been trained. In some embodiments visual data is being generated and transmitted across the network substantially simultaneously, for example during a live broadcast. In such embodiments, there may not be enough time to generate a specific hierarchical

algorithm for the visual data without introducing a substantial delay to the live broadcast. In such embodiments, selecting a hierarchical algorithm that has been pre-trained on similar visual data can reduce this delay.

Once hierarchical algorithms have been trained in order to enhance lower-quality versions of the new visual data these can be stored in order to be sent with the visual data or distributed to receiving nodes in advance of visual data being sent, depending on the embodiments.

These embodiments can be used in combination with other embodiments, and alternative and optional portions of embodiments, described elsewhere in this specification.

Strided Convolutions

Optionally, FIG. 18 shows an alternative method 1800 of the method 1500 shown in FIG. 15 according to some embodiments. In these embodiments, once the visual data 1510 has been separated into individual frames at step 1520, the dimension of the full resolution images 2110 extracted may be reduced based on at least one predetermined factor at step 1810. In these embodiments, at step 1810, the at least one predetermined factor may be used to select individual pixels from the extracted full resolution images 2110 to form a lower resolution representation 2120 of the extracted full resolution images 2110. For example, in some embodiments, a predetermined factor of 2 may indicate every other pixel in both the horizontal and vertical dimensions is selected, as shown in FIG. 21. It will be appreciated that other values of the predetermined factor may be used in other embodiments, furthermore in some of these other embodiments both the horizontal and vertical dimensions may each have a different predetermined factor applied.

In some embodiments, the reduced dimension visual data may be concatenated in step 1820, to form a lower resolution representation 2120 of the extracted full resolution images 2110. In these embodiments, the lower resolution representation 2120 can then be used as a basis for any image enhancement techniques and analysis allowing for the use of larger super resolution networks without compromising the runtime. In some embodiments, different predetermined factors and sizes of super resolution network can be used to obtain optimum performance.

Alternatively, in other embodiments, FIG. 19 shows a method 1900 based upon the method 1600 shown in FIG. 16. In method 1900 of these other embodiments, steps 1810 and 1820 can occur after the frames are grouped into scenes in step 1610.

Referring now to FIG. 20, a modified process 2000 according to some embodiments (based on process 1900 but which could equally be a modified version of process 1800) is shown.

In these embodiments, process 1900 is modified with further steps. Specifically, in some of these embodiments, step 1710 can begin with a high resolution visual data, for example at HD resolution. Then, at step 1720 of these embodiments, the visual data can be down-sampled to, for example, SD resolution to enable transmission over a bandwidth-restricted network. At step 1730 of these embodiments, the lower-resolution visual data can be transmitted over a network thus becoming the input visual data 310 of either process 1800 or 1900 described above in relation to FIG. 18 or 19 respectively.

Specific Description—Flexibility of Described Aspects and Embodiments

Any system feature as described herein may also be provided as a method feature, and vice versa.

As used herein, means plus function features may be expressed alternatively in terms of their corresponding structure.

In particular, method aspects may be applied to system aspects, and vice versa.

Furthermore, any, some and/or all features in one aspect can be applied to any, some and/or all features in any other aspect, in any appropriate combination.

It should also be appreciated that particular combinations of the various features described and defined in any aspects of disclosed implementations can be implemented and/or supplied and/or used independently.

In alternative embodiments, the input visual data concerned may be media for playback, such as recorded visual data or live streamed visual data, or it can be videoconference video or any other visual data source such as video recorded or being recorded on a portable device such as a mobile phone or a video recording device such as a video camera or surveillance camera.

It should also be appreciated that the term ‘visual data’, may refer to a single image, a sequence of images, video, or a section of visual data.

It should further be appreciated that the term “enhancing” may refer to upscaling, increasing the resolution and/or quality of visual data. References to enhancing or increasing the quality of visual data can refer to upscaling or using enhancement techniques of the possible embodiments described. References to down sampling can refer to reducing the resolution and/or quality of visual data (for example by quantisation to lower the bit rate of the visual data).

It should also be appreciated that the term ‘frame’, particularly in reference to grouping multiple frames into scenes, can refer to both an entire frame of a video and an area comprising a smaller section of a frame.

In aspects and/or embodiments, the terms algorithms and/or models and/or parameters can be used interchangeably or exchanged with each other. Further, in aspects and/or embodiments, the terms hierarchical algorithm, hierarchical model and hierarchical parameter can be exchanged with the terms convolutional neural networks and/or convolutional neural network model, convolutional neural network algorithm, convolutional neural network parameter.

What is claimed is:

1. A method for training learned hierarchical algorithms, the method comprising the steps of:

receiving input data;

generating metric data from the input data, the metric data measuring quality of output data produced for the input data by a plurality of pre-trained hierarchical algorithms stored in a library, wherein each of the plurality of pre-trained hierarchical algorithms is associated with respective metric data and each is trained on different input data;

selecting at least one hierarchical algorithm from the plurality of pre-trained hierarchical algorithms based on comparing the respective metric data;

training, using a deep learning approach, the at least one hierarchical algorithm based on the input data to generate a new trained hierarchical algorithm; and
adding the new trained a hierarchical algorithm to the library.

2. The method of claim 1, wherein each of the plurality of pre-trained hierarchical algorithms are stored with features extracted from the input data used in training the hierarchical algorithm.

3. The method of claim 1, wherein training the at least one hierarchical algorithm includes training the at least one hierarchical algorithm for a predetermined number of iterations.

4. The method of claim 1, wherein the metric data includes an error rate, a peak signal-to-noise ratio, a distance metric, a similarity measure, or a structural similarity index.

5. The method of claim 1, wherein training the at least one hierarchical algorithm includes training the at least one hierarchical algorithm until a predetermined value of reconstruction error is reached.

6. The method of claim 1, wherein the input data is a section of video and the metric data is based on a style of the section of video.

7. The method of claim 1, wherein the hierarchical algorithm is a non-linear hierarchical algorithm.

8. The method of claim 1, wherein the hierarchical algorithm performs image enhancement, using super-resolution techniques.

9. The method of claim 1, wherein the hierarchical algorithm uses a spatio-temporal approach.

10. The method of claim 1, wherein each of the pre-trained hierarchical algorithms in the library has a library reference identifier.

11. A method for training neural networks, the method comprising the steps of:

receiving input data;

generating metric data from the input data, the metric data measuring quality of output data produced for the input data by a plurality of pre-trained neural networks stored in a library, wherein each of the plurality of pre-trained neural networks is associated with respective metric data and each is trained on different input data;

selecting at least one neural network from the plurality of pre-trained neural networks based on comparing the respective metric data;

training, using a deep learning approach, the at least one neural network based on the input data to generate a new trained neural network; and

adding the new trained neural network to the library.

12. The method of claim 11, further comprising the step of:

storing the new trained neural network with features generated from the input data.

13. The method of claim 11, wherein the metric data is based on generating the highest quality output data, and wherein the quality is defined using an error rate, a peak signal-to-noise ratio, a distance metric, a similarity measure, or a structural similarity index.

14. The method of claim 11, wherein the at least one neural network is a convolutional neural network.

15. The method of claim 11, wherein the metric data is a similarity measure, generating the metric data from the input data includes extracting features of the input data and selecting the at least one neural network includes:

using the extracted features to select the at least one pre-trained neural network from the plurality of neural networks as associated with features most similar to the features of the input data.

16. The method of claim 11, wherein each of the pre-trained neural networks in the library has a library reference identifier.

17. The method of claim 11, wherein the respective metric data measures how optimal the pre-trained neural network is for providing initialisation data for generating the new neural network.

18. The method of claim **11**, wherein the input data comprises any of: a sequence of frames of visual data and a region within a sequence of frames of visual data.

19. The method of claim **11**, wherein the training occurs for a pre-determined number of iterations to limit computational time. 5

20. A computer program product embodied on a non-transitory storage medium and comprising instructions that, when executed, cause a system to train learned hierarchical algorithms, by performing the steps of: 10

receiving input data;

generating metric data from the input data, the metric data measuring quality of output data produced for the input data by a plurality of pre-trained hierarchical algorithms stored in a library, wherein each of the plurality 15 of pre-trained hierarchical algorithms is associated with respective metric data and each is trained on different input data;

selecting at least one hierarchical algorithm from the plurality of pre-trained hierarchical algorithms based 20 on comparing the respective metric data;

training, using a deep learning approach, the at least one hierarchical algorithm based on the input data to generate a new trained hierarchical algorithm; and

adding the new trained hierarchical algorithm to the 25 library.

* * * * *